Conditional Excess Volatility*

Christian Skov Jensen and Eben Lazarus

JULY 2026

Abstract

We decompose market return variance into a component that covaries with the stochastic discount factor (SDF) and one that does not. The SDF-orthogonal component — *conditional excess volatility* — contributes to variance but does not earn a risk premium. Time variation in the ratio of the market risk premium to total market variance reveals this unpriced share of market variance. Using three distinct measures of this ratio across twenty equity markets, we find that it does not increase during recessions and often declines. Because standard models predict countercyclical risk pricing, this finding requires the excess volatility share to increase substantially in bad times. Our empirical results imply that excess volatility constitutes at least half of the total market variance in recessions. We show that this mechanism provides a unified explanation for (i) why optimal portfolio weights do not increase in downturns despite rising Sharpe ratios, (ii) why the term structure of Sharpe ratios on dividend claims is downward-sloping, and (iii) why option portfolios with fixed risk exposures earn lower returns at longer horizons. Furthermore, high conditional excess volatility predicts lower subsequent investment, implying that unpriced fluctuations in aggregate market risk have important real effects.

*We are grateful to Mike Chernov (discussant), Niels Gormsen, Theis Ingerslev Jensen, Frank de Jong, Martin Lettau, Ian Martin, David Schreindorfer, Tobias Sichert, and seminar participants at Bocconi University, University of Copenhagen, Tilburg University, the European Summer Symposium in Financial Markets 2025, the 11th ESSEC Empirical Finance Workshop, and the Virtual Derivatives Seminar for very helpful comments. We thank Theis Ingerslev Jensen for sharing global returns data. Main parts of this paper were previously circulated under the title "The Cyclicality of Risk and Risk Premia." Jensen is at University of Copenhagen, Bocconi University, and the Danish Finance Institute (christian.jensen@unibocconi.it, csj@econ.ku.dk), and Lazarus is at the Haas School of Business, University of California, Berkeley, and NBER (lazarus@berkeley.edu).

1. Introduction

Market volatility is widely used as a real-time measure of aggregate risk. Standard asset-pricing theory implies, however, that not all return variation is priced. Under no arbitrage, the stochastic discount factor (SDF) determines risk premia through its covariance with returns, so variation in returns that is orthogonal to the SDF cannot command a risk premium. Market volatility therefore combines two conceptually distinct components: priced risk, which is compensated through expected returns, and unpriced fluctuations, which are not. We refer to the unpriced component as "conditional excess volatility": the share of market return variance that is irrelevant for marginal utility.

Conditional excess volatility arises whenever market variance becomes sufficiently large relative to plausible expected returns. Consider the Global Financial Crisis. As reported in Table 1, at the height of the crisis in 2008–2009, annualized market variance rose to roughly 69%. Even under the lowest plausible level of constant risk aversion, $\widetilde{\gamma} = 1$, such variance would imply an annualized risk premium on the order of 69%, far exceeding any reasonable estimate of expected returns at the trough. The gap between the risk premium that such variance would imply and the premium that can plausibly be expected reveals a component of return variation that is orthogonal to marginal utility. This is the logic of our measurement approach.

Our framework builds on a simple identification result. Let $\gamma_t \equiv \mu_t / \sigma_t^2$ denote the ratio of the conditional market risk premium to conditional market variance. We show that this ratio decomposes as

$$\gamma_t = \widetilde{\gamma}_t (1 - \psi_t), \tag{1}$$

where $\widetilde{\gamma}_t$ is the price of spanned SDF-risk and ψ_t is the share of market variance that is orthogonal to the SDF. The ratio γ_t therefore identifies the fraction of volatility that is relevant for marginal utility: departures of γ_t from $\widetilde{\gamma}_t$ directly measure the excess volatility share ψ_t.

This decomposition delivers clear cyclical implications. In standard asset-pricing models, $\widetilde{\gamma}_t$ is countercyclical, rising in bad times. If γ_t were constant, the decomposition

Table 1
Variance swings and implied risk premia

Panel A reports the twelve month pre-crisis annualized variance and the peak variance at the trough of recessions. We also report the ratio of the two variances. Panel B reports the annualized implied risk-premium ($\mu = \widetilde{\gamma}\sigma^2$) in crisis periods under constant risk aversion, $\widetilde{\gamma} = 1$ or 3.

Panel A: US Peak-to-Trough Variance			
Episode	Pre-crisis variance (ann.)	Trough variance (ann.)	Variance Ratio
Covid 2020	2.37%	98.61%	41.54x
GFC 2007-2009	2.21%	69.15%	31.22x
Dot-com 2001	5.26%	13.59%	2.58x
Panel B: Implied Crisis Risk Premium Under Constant $\widetilde{\gamma}$			
Episode	Variance Ratio	Implied μ if $\widetilde{\gamma} = 1$	Implied μ if $\widetilde{\gamma} = 3$
Covid 2020	41.54x	98.61%	295.82%
GFC 2007-2009	31.22x	69.15%	207.46%
Dot-com 2001	2.58x	13.59%	40.78%

would mechanically imply that ψ_t must increase during downturns: as the price of SDF-risk rises, a larger share of market variance would have to become orthogonal to the SDF in order to keep γ_t unchanged. In the data, however, we find that γ_t is weakly procyclical. This procyclicality implies that ψ_t must be even more countercyclical than under a constant γ_t, increasing precisely in recessions and crises when market volatility typically spikes. Recessions are therefore characterized by a systematic shift in the composition of aggregate risk toward its unpriced component. Figure 1 illustrates this pattern, showing the implied share of aggregate market volatility that is unpriced, along with the decomposition into priced and unpriced components, under the conservative assumption that the price of spanned risk is constant over time.[1]

Our decomposition admits several complementary interpretations that clarify why γ_t is an informative statistic for investigating excess volatility. First, γ_t equals the conditional beta of the SDF with respect to the market return (up to scaling by the risk-free rate). Excess volatility dilutes this beta: when ψ_t is high, a unit of market variance contains less

[1] Two notes on the figure. First, conditional excess volatility is computed using the expected measure discussed in Section 3 and is capped at 0.9. Second, if total market volatility were priced, the implied average excess return would be average$(\sigma_t^2 \times \tilde{\gamma}) = 21.72\%$ annually, much higher than the historical average excess return of 7.23% in this sample. Using only the priced aggregate risk, the average excess return is 5.91%, much closer to the historical average.

Figure 1. Priced and unpriced market volatility. Panel A shows the annualized monthly realized volatility for the US stock market. The shaded dark gray area represents conditional excess volatility. Panel B shows the share of market volatility that is unpriced. In this illustration, we set the price of spanned risk $\tilde{\gamma} = 8$. Vertical shaded gray areas are NBER recession periods.

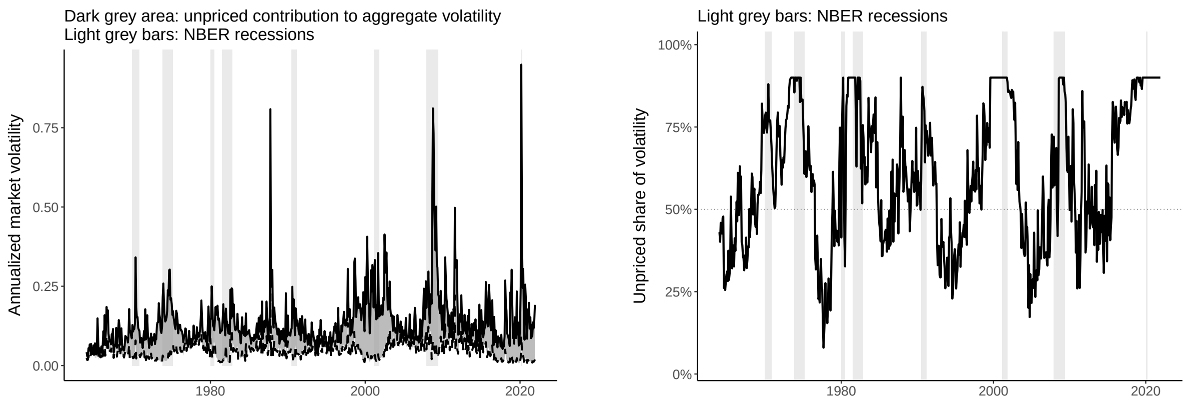

Panel A. Volatility decomposition   Panel B. Conditional excess volatility

SDF-relevant information, implying that the SDF loads less on the market. Second, γ_t determines optimal portfolio choice. Under the standard myopic portfolio rule, an investor with risk aversion $\widetilde{\gamma}$ allocates $w_m = \gamma_t / \widetilde{\gamma}$ to the market, implying $w_m = 1 - \psi_t$ when risk aversion is constant; higher excess volatility in bad times therefore leads investors to optimally reduce market exposure in bad times. Third, γ_t is often interpreted as a measure of representative-agent risk aversion, but our decomposition shows this interpretation breaks down when ψ_t is large because unpriced variance inflates the denominator of μ_t/σ_t^2. Fourth, the same logic extends to the term structure: long-maturity claims are more exposed to discount-rate shocks, implying they have higher $\psi_{n,t}$ and lower $\gamma_{n,t}$ than short-maturity claims. While γ_t and the Sharpe ratio of market returns are tightly linked ($\gamma_t = \text{Sharpe ratio}_t / \sigma_t$), these interpretations are unavailable for the Sharpe ratio.

Because the excess volatility share ψ_t is not directly observable, we study its cyclicality through the observable ratio γ_t. We estimate γ_t using three complementary approaches that rely on distinct sources of information. First, we construct a realized measure ($\gamma_t^{\text{realized}}$) using within-month daily returns, which provides an ex post estimate based on realized excess returns and realized variance. Second, we compute an expected measure ($\gamma_t^{\text{expected}}$) that combines ex ante predictions of the market risk premium from Kelly and Pruitt (2013), based on the cross-section of valuation ratios, with forecasts of market variance from a simple autoregressive model. Third, we infer an option-implied measure (γ_t^{option}) from option prices by estimating the projection of the SDF onto market returns. These three measures — realized, expected, and option-implied — provide independent estimates of the price per unit of variance risk, γ_t. Despite their different foundations, the three measures exhibit similar time-series behavior, with pairwise correlations between 0.27 and 0.42.

Across the three approaches, we find that γ_t does not increase in recessions; in the modern sample starting in 1964, it typically decreases. Interpreted through the decomposition $\gamma_t = \widetilde{\gamma}_t(1 - \psi_t)$, this pattern implies that the excess volatility share ψ_t increases in bad times. Even though the individual estimates of γ_t are measured with statistical noise, the implication is nonetheless strong. Standard asset-pricing models predict that the price of SDF-risk $\widetilde{\gamma}_t$ is countercyclical, implying that γ_t should also increase when economic conditions deteriorate. A failure of γ_t to increase therefore already requires a substantial increase in the share of market variance that is orthogonal to the SDF. The fact that γ_t instead decreases in recessions strengthens this conclusion further: a growing fraction of market volatility in downturns must reflect unpriced, or excess, volatility. Consistent with this interpretation, an empirical quantification of ψ_t indicates that the excess volatility share is at least 50% during recessions, implying that at least half of total market variance in downturns reflects excess variance rather than SDF-relevant variation.

Additional macroeconomic evidence further supports this interpretation. The estimated γ_t covaries strongly with standard indicators of the business cycle: it is procyclical, rising when recession probabilities are low and financial conditions are loose, and falling when recession risk increases and financial conditions tighten. Similarly, γ_t comoves positively with real economic activity, including consumption growth and industrial production growth. These patterns are difficult to reconcile with standard asset-pricing models, which predict that γ_t should increase in bad times as $\widetilde{\gamma}$ increases. Instead, the observed procyclicality implies that the excess volatility share ψ_t must increase during downturns, diluting the link between market variance and marginal utility precisely when volatility spikes. Importantly, this pattern is not unique to the United States. Using data from 20 international stock market indexes, we find similar procyclical behavior in γ_t across countries, suggesting that the rise in excess volatility during recessions is a pervasive feature of global equity markets rather than a country-specific phenomenon.

Our results are closely related to, but conceptually distinct from, the findings of Moreira and Muir (2017), who show that investors can earn alpha by mechanically scaling market exposure with realized variance. We instead measure the fraction of market variance that is priced, conditioning on business-cycle variables rather than on realized variance itself. Our framework then provides a structural interpretation of why their volatility-managed strategy works. When variance spikes, a disproportionate share of that volatility reflects excess variance — fluctuations that are orthogonal to the SDF and therefore do not command a risk premium. Scaling down exposure in high-volatility periods thus reduces the investor's exposure to uncompensated risk. In our framework, this mechanism appears through the decomposition $w_m = 1 - \psi_t$: when the excess volatility share ψ_t rises, the optimal portfolio weight declines even if the Sharpe ratio is high. Volatility-managed strategies earn alpha precisely because increases in market variance tend to coincide with increases in excess volatility.

The rise in unpriced volatility in downturns also has real consequences. High conditional excess volatility predicts lower investment growth over the following quarters, and this predictability is robust to controlling for standard uncertainty and business-cycle measures, including the VIX, recession probabilities, financial conditions, valuation ratios, and credit spreads. Because excess volatility does not command a risk premium by construction, this evidence indicates that firms respond to fluctuations in uncertainty that are not reflected in expected returns.

To illustrate the mechanism behind our empirical findings, we develop a stylized asset-pricing model with an explicit source of excess volatility. The model builds on Lettau and Wachter (2007) but augments it with unpriced, non-fundamental shocks to dividends. In the model, market return volatility therefore has three components: constant dividend (fundamental) volatility, constant discount-rate volatility, and non-fundamental volatility that increases in bad times. When the price of fundamental risk rises, non-fundamental return risk rises as well, increasing total market variance without a proportional increase in expected returns. This mechanism generates a procyclical γ_t, consistent with our empirical estimates, while preserving a countercyclical Sharpe ratio. It also produces a downward-sloping term structure of Sharpe ratios because long-maturity claims are more exposed to unpriced discount-rate shocks, which raise their volatility without raising expected returns. We calibrate and simulate the model and find that it matches the main empirical estimates both qualitatively and quantitatively.

As a further out-of-sample test, we exploit the fact that the procyclicality of γ_t implies a specific, testable prediction for option portfolios at different horizons: strategies that hold fixed the quantity of risk (binary bets over index return states) should earn lower expected returns at longer horizons if γ_t is procyclical. Using S&P 500 index options, we find a downward-sloping term structure of risk prices for such strategies, providing direct support for the procyclicality result of γ_t and the necessity for ψ_t to increase in bad times using an independent data source.

Related literature. Our paper relates to several strands of the asset-pricing literature. The classic excess volatility literature, initiated by Shiller (1981) and LeRoy and Porter (1981) and later extended by Campbell and Shiller (1987), shows that stock prices fluctuate more than can be justified by subsequent dividend realizations under standard present-value restrictions. Our paper differs from this literature by proposing a new benchmark for excess volatility that does not require specifying a model of cash flows or discount rates; instead, we use only the no-arbitrage condition to measure excess volatility as the component of return variance orthogonal to the SDF. Our benchmark is conditional by construction and allows us to investigate *when* excess volatility is high.

Our work is also related to studies of volatility-managed portfolios such as Moreira and Muir (2017), who show that scaling market exposure by realized variance generates significant alpha. We provide a structural interpretation for this result by showing that spikes in market variance disproportionately reflect excess, SDF-irrelevant volatility. Schreindorfer and Sichert (2023) study the cyclical behavior of the projection of the SDF onto the market and document patterns consistent with procyclical risk pricing.

Our analysis also connects to work on the term structure of equity risk premia. Lettau and Wachter (2007) develop a model that generates a downward-sloping term structure of Sharpe ratios for dividend strips, while Gormsen (2021) provides empirical evidence on the term structure of equity risk premia using dividend futures. We build on this framework by introducing an explicit source of unpriced, non-fundamental volatility that raises return variance in bad times without proportionally increasing expected returns.

Our option-implied measures relate to a longstanding literature on option-implied risk aversion (e.g., Ait-Sahalia and Lo, 2000; Jackwerth, 2000; Rosenberg and Engle, 2002; Bliss and Panigirtzoglou, 2004), and our option portfolio term structure analysis extends results in Bliss and Panigirtzoglou (2004) with weaker parametric assumptions.

The rest of the paper is structured as follows. In Section 2, we develop the concept of conditional excess volatility. Section 3 describes how we measure γ_t in the data. Section 4 presents empirical evidence on the cyclicality of γ_t, quantifies the required rise in excess volatility, and shows that high excess volatility predicts lower subsequent real investment. Section 5 investigates excess volatility under specific parametric assumptions about the SDF. In Section 6, we develop a model with excess volatility and show simulation and quantification results of the model. Section 7 presents further evidence on the cyclicality of γ_t and ψ_t from the term structure of specific option portfolios. Section 8 concludes the paper.

2. Conditional Excess Volatility

In this section, we show that the absence of arbitrage implies a natural decomposition of market return variance into a component spanned by the SDF and an orthogonal residual. The residual — return variation that does not covary with marginal utility — is conditional excess volatility. We show that γ_t identifies the time-varying share of excess volatility, that this share must increase in crises, and that the decomposition connects to portfolio choice and the equity term structure.

2.1. The no-arbitrage decomposition

Assuming the absence of arbitrage, there exists a strictly positive one-period SDF M_{t+1} such that $\mathbb{E}_t[M_{t+1}R_{t+1}] = 1$ for any gross return R_{t+1}. This implies that the risk premium on any asset is determined entirely by its covariance with M:

$$\mu_t \equiv \mathbb{E}_t[R_{m,t+1} - R_{f,t+1}] = -R_{f,t+1}\text{Cov}_t(M_{t+1}, R_{m,t+1}). \tag{2}$$

Nothing else about the return distribution — in particular, not the total variance $\sigma_t^2 \equiv \text{Var}_t(R_{m,t+1})$ — enters into risk premia except through this covariance. Because only $\text{Cov}_t(M_{t+1}, R_{m,t+1})$ matters for pricing, we can project the market return onto the SDF:

$$R_{m,t+1} = a_t + \beta_{R\to M,t}\, M_{t+1} + u_{t+1}, \qquad \text{Cov}_t(u_{t+1}, M_{t+1}) = 0, \tag{3}$$

where $\beta_{R\to M,t} \equiv \text{Cov}_t(R_{m,t+1}, M_{t+1})/\text{Var}_t(M_{t+1})$ is the slope coefficient in a conditional linear projection and u_{t+1} is the projection residual. This decomposition is uniquely determined by the joint distribution of $(R_{m,t+1}, M_{t+1})$.

Total variance decomposes as:

$$\sigma_t^2 = \underbrace{\beta_{R\to M,t}^2 \cdot \text{Var}_t(M_{t+1})}_{\equiv \sigma_{S,t}^2 \text{ (spanned)}} + \underbrace{\text{Var}_t(u_{t+1})}_{\equiv \sigma_{\perp,t}^2 \text{ (excess)}}. \tag{4}$$

The spanned variance $\sigma_{S,t}^2$ is the component of return variance attributable to the SDF. The residual $\sigma_{\perp,t}^2$ is return variance orthogonal to the SDF. From equation (2), the risk premium depends only on the spanned component:

$$\mu_t = -R_{f,t+1}\, \beta_{R\to M,t} \cdot \text{Var}_t(M_{t+1}). \tag{5}$$

The variance that is orthogonal to the SDF, $\sigma_{\perp,t}^2$, has no effect on the risk premium. This is a direct implication of equation (2): variance that does not covary with M cannot generate expected returns.

We refer to $\sigma_{\perp,t}^2$ as *conditional excess volatility* — return variation that, by the fundamental theorem of asset pricing, is not compensated by a risk premium.[2] In an economy where

[2]The term "excess volatility" connects to Shiller (1981) and LeRoy and Porter (1981), who documented that stock prices are more volatile than the present value of subsequent dividends. The classical notion is unconditional and fundamentals-based; ours is conditional and SDF-based. We return to this connection in Section 2.4.

the market perfectly spans the SDF — as, for example, in the CAPM — $\sigma^2_{\perp,t} = 0$ and all return variance commands a risk premium. In any economy with risks beyond the market (labor income, housing, private equity, or non-fundamental sources of return variation), the SDF has components not captured by $R_{m,t+1}$, and generically $\sigma^2_{\perp,t} > 0$.

2.2. Identification through γ_t

We now show that γ_t identifies the share of total variance that is excess volatility. Define the *excess volatility share*

$$\psi_t \equiv \frac{\sigma^2_{\perp,t}}{\sigma^2_t} \in [0, 1), \tag{6}$$

and the *price per unit of spanned variance*

$$\widetilde{\gamma}_t \equiv \frac{\mu_t}{\sigma^2_{S,t}}. \tag{7}$$

The quantity $\widetilde{\gamma}_t$ is the risk premium per unit of variance that actually covaries with the SDF. The observable ratio $\gamma_t = \mu_t / \sigma^2_t$ divides the same risk premium by *total* variance, including the uncompensated component.[3] Their relationship is as follows.

Proposition 1 (Excess Volatility and the Price of Variance Risk)**.** *Under no-arbitrage:*

$$\gamma_t = \widetilde{\gamma}_t \cdot (1 - \psi_t). \tag{8}$$

Proof. $\gamma_t = \frac{\mu_t}{\sigma^2_t} = \frac{\mu_t}{\sigma^2_{S,t}} \cdot \frac{\sigma^2_{S,t}}{\sigma^2_t} = \widetilde{\gamma}_t \cdot (1 - \psi_t).$ □

Equation (8) has three immediate implications. First, if the market perfectly spans the SDF, then $\psi_t = 0$ and $\gamma_t = \widetilde{\gamma}_t$: dividing by total variance and by spanned variance gives the same answer. Second, with excess volatility ($\psi_t > 0$), γ_t understates the price per unit of spanned variance: some of the variance in the denominator of γ_t is uncompensated. Third, time variation in γ_t reflects time variation in either $\widetilde{\gamma}_t$ or ψ_t or both. Under the maintained assumption that $\widetilde{\gamma}_t$ does not decrease in crises — shared by all standard asset-pricing

[3] The ratio γ_t has a long history: Fama (1968) refers to it as the "market price per unit of risk," and Friend and Blume (1975) use a version of it to measure aggregate relative risk aversion.

models — the procyclicality of γ_t requires ψ_t to increase in bad times.

Corollary 1 (Constant Risk Aversion)**.** *If $\widetilde{\gamma}_t = \widetilde{\gamma}$ is constant, then $\psi_t = 1 - \gamma_t/\widetilde{\gamma}$, and all time variation in γ_t is driven by time variation in the excess volatility share.*

2.3. Why excess volatility must increase in crises

The decomposition in Proposition 1 allows us to make a sharp argument about crisis episodes. During the 2008–2009 financial crisis, annualized market variance reached roughly 69%. Under *any* constant risk aversion $\widetilde{\gamma}$ — even the lowest plausible value, $\widetilde{\gamma} = 1$ — the identity $\mu_t = \widetilde{\gamma}\,\sigma_t^2$ would imply an annualized risk premium of 69%, far exceeding any plausible estimate of expected excess returns at the trough. Under a more standard calibration $\widetilde{\gamma} = 3$, the implied risk premium would be 207%. The gap between the risk premium that this variance would imply and the premium it actually generates reveals conditional excess volatility.

The following proposition formalizes this argument.

Proposition 2 (Lower Bound on Crisis Excess Volatility)**.** *Let G and B be index expansion and crisis states, respectively. Suppose $\widetilde{\gamma}_B \geq \widetilde{\gamma}_G$ (the price per unit of spanned variance does not decrease in crises). Then:*

$$\psi_B \geq 1 - \frac{\gamma_B}{\gamma_G}\,(1 - \psi_G). \tag{9}$$

If in addition $\psi_G = 0$ (no excess volatility in good times, the most conservative assumption):

$$\psi_B \geq 1 - \frac{\gamma_B}{\gamma_G}. \tag{10}$$

Proof. From Proposition 1: $\gamma_B = \widetilde{\gamma}_B(1 - \psi_B)$ and $\gamma_G = \widetilde{\gamma}_G(1 - \psi_G)$. Since $\widetilde{\gamma}_B \geq \widetilde{\gamma}_G$:

$$\frac{\gamma_B}{\gamma_G} = \frac{\widetilde{\gamma}_B}{\widetilde{\gamma}_G} \cdot \frac{1 - \psi_B}{1 - \psi_G} \geq \frac{1 - \psi_B}{1 - \psi_G}.$$

Rearranging: $1 - \psi_B \leq (\gamma_B/\gamma_G)(1 - \psi_G)$. □

Example. Our empirical estimates in Section 4 show that γ_t roughly halves in crises relative to expansions (using either the expected or option-implied measures in the modern sample). If $\gamma_G \approx 2\gamma_B$, then under zero excess volatility in good times, equation (10) implies $\psi_B \geq 0.50$: at least half of crisis-time market variance is excess volatility. This is a lower bound — it assumes no excess volatility in good times and that the price per unit of spanned variance does not increase in crises (the weakest possible assumptions). Under more realistic assumptions, the excess volatility share would be higher.

A natural response to the finding that γ_t is only weakly procyclical is to question its economic significance. The next proposition shows that even modest procyclicality — or mere acyclicality — requires large swings in excess volatility.

Proposition 3 (Amplification)**.** *Suppose the price per unit of spanned variance is countercyclical:* $\widetilde{\gamma}_B = k \cdot \widetilde{\gamma}_G$ *for some* $k > 1$*. If* $\gamma_B = \gamma_G$ *(i.e.,* γ_t *is merely flat, not even procyclical), then:*

$$\psi_B = 1 - \frac{1 - \psi_G}{k}. \tag{11}$$

Proof. $\gamma_G = \widetilde{\gamma}_G(1 - \psi_G)$ and $\gamma_B = k\widetilde{\gamma}_G(1 - \psi_B) = \gamma_G$ imply $(1 - \psi_B) = (1 - \psi_G)/k$. □

The implications of Proposition 3 are summarized in Table 2. Suppose fundamental risk aversion doubles in crises ($k = 2$), which is conservative for habit formation and long-run risk models. If the excess volatility share is 20% in good times, then merely flat γ_t requires it to triple in crises, to at least 60%. If γ_t actually *declines* in crises, as the data suggest, the required increase is larger.

Table 2
Required excess volatility for flat γ_t

This table reports the crisis-time excess volatility share ψ_B required for γ_t to be merely flat ($\gamma_B = \gamma_G$), given different assumptions about the increase in fundamental risk aversion ($k = \widetilde{\gamma}_B / \widetilde{\gamma}_G$) and the good-times excess volatility share (ψ_G).

k	ψ_G	Required ψ_B	$\Delta\psi$
1.5	0.00	0.33	+0.33
1.5	0.20	0.47	+0.27
2.0	0.00	0.50	+0.50
2.0	0.20	0.60	+0.40

Put differently, the procyclicality of γ_t is weak only in a statistical sense (the estimates are noisy). The economic implication is strong: even the most conservative reading of the data (that γ_t does not rise in recessions) is incompatible with standard models unless excess volatility increases substantially in bad times. And the relevant benchmark is the *countercyclical* γ_t predicted by every standard model, rather than a constant γ_t.

2.4. Connection to classical excess volatility tests

Shiller (1981) showed that stock prices are more volatile than the present value of subsequent dividends, launching a large literature on excess volatility (LeRoy and Porter, 1981). Any such test must take a stand on the benchmark level of volatility that fundamentals justify.

Our notion of excess volatility is related to but distinct from Shiller's, as highlighted in Table 3. His benchmark is the present value of realized dividends under a given discount rate model; violations indicate that discount rates vary or that markets are inefficient. We instead define excess volatility directly from the SDF: $\sigma^2_{\perp,t}$ is the variance of the component of $R_{m,t+1}$ orthogonal to M_{t+1}. No model for cash flows or discount rates is required — only the absence of arbitrage.

Table 3
A new perspective on excess volatility

	Shiller (1981)	This paper
Criterion	Cash-flow fundamentals	Covariance with the SDF
Type	Unconditional	Conditional (time-varying)
Requires	Discount rate model	No-arbitrage only
Connects to portfolio choice	No	Yes ($w_m = 1 - \psi_t$)
Connects to term structure	Indirectly	Directly

While the two notions can overlap, they need not do so. For example, suppose discount-rate variation is rational and that the market SDF fully prices those fluctuations. In such an economy, prices move too much relative to dividends but all the variance may still be SDF-related. Consequently, the classical notion of excess volatility exists but the SDF-orthogonal excess volatility is zero. This is essentially the modern rational-discount-rate interpretation of Shiller's result.

Alternatively, suppose prices satisfy present-value relations but large volatility spikes occur and these events are unpriced. Then there may be no violation of dividend-based valuation yet variance may be excess in the SDF sense. These examples illustrate that while the two notions are related, they are also conceptually distinct.

Under both notions, markets fluctuate more than compensation for risk requires. Shiller's test is whether they fluctuate more than what dividends justify. Ours is instead whether they fluctuate more than justified given changes in expected returns. Our notion is more general in the sense that it applies to any no-arbitrage economy, and it is specifically designed to be conditional — allowing us to ask *when* excess volatility is large.

2.5. Economic implications

SDF beta interpretation. The ratio γ_t equals the conditional beta of the SDF on the market return, up to a scaling by $R_{f,t+1}$: $\gamma_t = -R_{f,t+1}\beta_{M\to R,t}$, where $\beta_{M\to R,t} = \text{Cov}_t(M_{t+1}, R_{m,t+1})/\sigma_t^2$. From Proposition 1, excess volatility *dilutes* this beta. When ψ_t is high, a unit of market variance contains less SDF-relevant information, and γ_t is correspondingly low. The procyclicality of γ_t therefore means that the SDF loads less heavily on the market in bad times, reflecting the rise in excess volatility rather than any decline in fundamental risk.

Portfolio choice. Under the standard myopic portfolio choice rule, an investor with relative risk aversion $\widetilde{\gamma}$ invests a share $w_m = \mu_t/(\widetilde{\gamma}\,\sigma_t^2) = \gamma_t/\widetilde{\gamma}$ of wealth in the market. From Proposition 1, if $\widetilde{\gamma}_t = \widetilde{\gamma}$:

$$w_m = 1 - \psi_t. \tag{12}$$

The optimal portfolio weight equals one minus the excess volatility share. When ψ_t is high (bad times), the investor allocates *less* to the market, even though the Sharpe ratio is elevated. The textbook prescription to maximize the Sharpe ratio would counsel moving into the market in bad times. But the Sharpe ratio treats all volatility symmetrically; the optimal weight correctly accounts for the fact that only spanned variance commands a risk premium.

This provides a precise explanation for the findings of Moreira and Muir (2017): volatility-managed strategies that scale down in high-variance regimes earn alpha because they avoid taking uncompensated risk. The strategy's alpha is proportional to the covariance between ψ_t and σ_t^2.

Risk aversion approximation. Under standard assumptions (log-normal returns, power utility), γ_t approximates the representative investor's relative risk aversion. From Proposition 1, this approximation breaks down whenever ψ_t is large. In crisis episodes — precisely when one might most want to measure risk aversion — γ_t understates $\widetilde{\gamma}_t$ because the denominator includes uncompensated variance. This bias is procyclical: it is smallest in good times (when ψ_t is low) and largest in bad times (when ψ_t is high).

Term structure. For a zero-coupon equity claim of maturity n, the analogous decomposition gives $\gamma_{n,t} = \widetilde{\gamma}_{n,t} \cdot (1 - \psi_{n,t})$. Long-maturity claims have greater exposure to discount-rate shocks, which in standard models are either unpriced or carry a small risk price relative to cash-flow risk. As a result, $\psi_{n,t}$ is increasing in n: short-maturity claims are mostly cash-flow risk (low ψ), while long-maturity claims carry more discount-rate risk (high ψ). This generates a downward-sloping term structure of $\gamma_{n,t}$ and, given the relationship $\text{SR}_{n,t} = \gamma_{n,t}\,\sigma_{n,t}$, a downward-sloping Sharpe ratio term structure — consistent with van Binsbergen and Koijen (2017) and Gormsen (2021). The model in Section 6 makes this mechanism explicit.

Real decisions. The distinction between priced and unpriced aggregate risk extends beyond financial markets. Firms make investment decisions under uncertainty, and households make consumption and labor-supply decisions in environments where future outcomes are uncertain. These decisions need not depend exclusively on compensated risk. Consequently, periods characterized by high levels of unpriced aggregate risk may have important real effects even though they generate little additional compensation in financial markets. A rise in unpriced aggregate risk increases uncertainty surrounding future outcomes without necessarily increasing expected returns, potentially encouraging firms to delay irreversible investment and households to postpone major expenditures. In the empirical analysis below, we show that measures of unpriced aggregate risk predict lower subsequent investment even after controlling for conventional measures of uncertainty, financial conditions, and business-cycle fluctuations.

3. Measuring γ_t

The excess volatility share (ψ_t) is not directly observable. However, as shown in Propositions 1 and 2, we can make statements about required time variation in the excess volatility from cyclical patterns in γ_t. In this section, we describe how we measure γ_t using three fundamentally different approaches.

Realized conditional price per unit of variance risk — $\gamma_t^{\textbf{realized}}$. Our first measure of the price per unit of variance risk is an ex-post measure that relies on realized within month daily returns. Let $\tilde{R}_s = R_s - R_s^f$ be the excess return on date s. We compute the realized conditional price per unit of variance risk in month t as

$$\gamma_t^{\text{realized}} = \frac{\sum_{s=1}^{N_t} \tilde{R}_s}{\frac{N_t}{N_t - 1} \sum_{s=1}^{N_t} [\tilde{R}_s - (\sum_{s=1}^{N_t} \tilde{R}_s)]^2} \tag{13}$$

where N_t denotes the number of trading days in month t.[4]

Expected conditional price per unit of variance risk — $\gamma_t^{\textbf{expected}}$. Our second approach of computing the price per unit of variance risk relies on ex ante predicted values for the conditional market risk premium and its variance. We compute the conditional market risk premium using the methodology in Kelly and Pruitt (2013). Under the two assumptions: (i) the expected log return and log growth rates are linear in a set of latent factors, and (ii) these factors evolve according to a first-order vector autoregression, Kelly and Pruitt (2013) show how to infer the conditional market risk premium from the cross-section of valuation ratios. The main reason why we choose this estimator as our ex ante predictor of the market risk premium is that the estimator does well in predicting market returns both in- and out-of-sample. Kelly and Pruitt (2013) find that they can predict the one-month market risk premium on the U.S. market portfolio with an R^2 of 2.38% in-sample and an R^2 of 0.93% out-of-sample. Gormsen and Jensen (2024) extend these results to 20 international stock market indexes, showing that the predictability exists everywhere. A minor benefit

[4]We thank Theis Ingerslev Jensen for sharing daily excess returns on international stock market indexes.

is that the method is built around valuation ratios and the predicted expected returns are therefore likely to fluctuate with the business cycle as we would expect.

Consistent with previous literature that considers time-variation in market variance (see e.g. Campbell et al. (2018)), we compute conditional expected market variance assuming that the variance follows a first-order autoregressive process. We compute the predicted variance via the relationship

$$\tilde{\text{Var}}_t(R_{m,t+1}) = \theta_0 + \theta_1 \text{Var}_{t-1}(R_{m,t}) \tag{14}$$

where

$$\text{Var}_{t-1}(R_{m,t}) = \frac{N_t}{N_t - 1} \sum_{s=1}^{N_t} [\tilde{R}_s - (\sum_{s=1}^{N_t} \tilde{R}_s)]^2 \tag{15}$$

is the realized variance in month t. We infer the values of the parameters θ_0 and θ_1 in equation (14) from a linear regression of realized variance on its one period lagged value.

Combining the predictions from Kelly and Pruitt (2013) about the conditional market risk-premium with the AR(1) variance prediction in equation (14), we compute $\gamma_t^{\text{expected}}$ as the ratio of the predicted risk premium to the predicted variance.

Option implied conditional price per unit of variance risk — γ_t^{option}. As our third approach for computing the price per unit of variance risk, we look to option markets.[5] Our method builds on and extends results in Bliss and Panigirtzoglou (2004). Imposing a functional form on the projection of the SDF onto the market return, $M_{t+1}|R_{m,t+1} = \delta_t R_{m,t+1}^{\gamma_t}$, we infer parameters δ_t and γ_t that best reconcile the ex ante observable option prices with the ex post realized returns. While Bliss and Panigirtzoglou (2004) impose that γ_t is time-invariant and let δ_t be free, we allow for γ_t to be time-varying while adding structure to δ_t. With this alternative parameterization, we rely on the so-called Berkowitz test, cf. Berkowitz (2001) together with constraints in Jensen, Lando, and Pedersen (2019) to infer the optimal time-varying sequence of risk prices, γ_t. For a detailed description of the method, see Appendix A.1.

[5]The empirical methodology in this section was first reported in Chapter 3 of Jensen (2018), which this paper now supersedes.

The Sharpe ratios. We compute the Sharpe ratios by multiplying the estimated price per unit of variance risks by the conditional volatility. For the realized Sharpe ratio, we multiply $\gamma_t^{realized}$ with the within month standard deviation of returns. For the expected Sharpe ratio, we multiply $\gamma_t^{expected}$ with the expected volatility from the AR(1) model. Lastly, for the option Sharpe ratio, we multiply γ_t^{option} with the option implied volatility computed from equation (45) using γ_t^{option} as the exponent.

Pairwise correlations. Panel A of Figure 2 shows the estimated Sharpe ratio and Panel B shows the price per unit of variance risk for the US stock market. As seen from the figure and reported in Table 4, the estimates for the prices of risks are highly correlated with correlations ranging from 0.27 to 0.42. In the table, we also report correlations for γ_t^{PCA} and SR_t^{PCA}. These are the first principal components for the price per unit of variance risk measures and the Sharpe ratio measures respectively. The first principal components capture a large part of the variation in the three measures with individual correlations to the measures ranging from 0.41 to 0.80. These principal components suggest that there is a strong similarity in the variation in the measures.

4. Excess Volatility over the Business Cycle

Given the three measures computed in the previous section, we now seek to investigate required variation in excess volatility over the business cycle. For fixed $\widetilde{\gamma}_t$, Proposition 1 shows that the cyclicality of γ_t directly identifies the cyclicality of excess volatility. In this section, we therefore study the cyclical behavior of excess volatility through the empirical cyclicality of γ_t.

4.1. Evidence from recessions

A natural starting point is to ask whether γ_t rises in recessions. However, as emphasized by Lustig and Verdelhan (2012), recessions are defined ex post as periods of significant declines in economic activity, typically accompanied by contractions in industrial production, employment, and GDP. In many cases — especially when recessions originate in financial markets, such as during the Global Financial Crisis — large declines in stock prices occur early in the downturn. Comparing the price of risk in recessions and non-recessions therefore risks introducing an ex post classification bias, because the timing of recessions is partly determined by the very market outcomes we are studying.

To address this issue, we adopt the following narrative benchmark. Consider an investor who wishes to enter the market during recessions and hold the market for twelve months thereafter. If the investor could select the entry point with perfect hindsight, a natural choice would be the market trough within each recession. We therefore consider a "perfect foresight" investor who can identify the market low in recessions in real time, enters the market at that point, and holds the market for the subsequent twelve months. This benchmark intentionally stacks the deck in favor of recession investing: if risk prices are countercyclical, such an investor should experience particularly attractive risk–return trade-offs relative to an investor who simply holds the market during normal times.

Panel A of Table 5 reports the difference in average monthly Sharpe ratios between the perfect-foresight recession investor and the normal-times investor. We compute this statistic for three samples reflecting the availability of the different γ_t measures: (i) the full sample starting in 1928, (ii) a post-1964 sample, and (iii) a post-1996 sample. The unconditional Sharpe ratios are calculated using all monthly excess returns earned during the twelve months following each recession trough. Consistent with the existing literature (e.g., Lustig and Verdelhan (2012)), we find that Sharpe ratios tend to be higher in recessions: the difference is positive and statistically significant in five of nine specifications.

Panel B turns to the behavior of the price of variance risk, γ_t. Here the evidence differs sharply from the Sharpe ratio patterns. While the long historical sample yields a positive and statistically significant difference, the modern samples — where expected and option-implied measures are available — show the opposite pattern. In these samples, γ_t is lower around recession troughs and the differences are statistically significant in several specifications. Taken together, the results indicate that the price of variance risk is at most acyclical and, if anything, procyclical in the modern data. Through the decomposition $\gamma_t = \widetilde{\gamma}_t(1 - \psi_t)$, this implies that the excess volatility share ψ_t rises in bad times.

To quantify the lower bound on how much excess volatility must increase in bad times,

Table 6 reports the minimum excess volatility share implied by the observed cyclicality in γ_t. From the decomposition in Propositions 2 and 3, the excess volatility ratio in bad times must satisfy

$$\psi_B = 1 - \frac{\gamma_B}{\gamma_G}\frac{1-\psi_G}{k}, \tag{16}$$

where γ_B and γ_G denote the price of variance risk in bad and good times, respectively. The table reports the ratio of risk prices, the minimum excess volatility share in good times ψ_G, and the implied minimum ψ_B. We set the relative increase in the price of SDF risk, $k = \widetilde{\gamma}_B/\widetilde{\gamma}_G = 1, 2$, or 4. A value of 1 corresponds to a constant price of spanned SDF-risk, a value of 2 to a doubling in downturns, and a value of 4 to a quadrupling of the SDF-relevant risk price. Even in specifications where γ_t rises slightly in recessions, the increase in risk aversion implies that excess volatility must still rise sharply in bad times. Across all specifications, the average lower bound on excess volatility is 0.09 in good times and 0.43 in bad times, reflecting that a much larger share of volatility in bad times is excess volatility that is orthogonal to the SDF.

The table also reports 90% bootstrapped confidence intervals for the ψ estimates. Across the 27 specifications, the lower bound for ψ_G exceeds zero in only six cases, and is statistically significant in just four of these, all corresponding to the highly conservative choice of $k = 1$. For any higher value of k, ψ_G is statistically indistinguishable from zero. In contrast, the results for ψ_B are markedly different. The point estimate is positive in 21 out of 27 specifications and statistically significant in 15 of these cases. Taken together, these findings provide strong evidence that the excess volatility share rises substantially during recessions. In the modern sample, and for the more realistic choice of $k = 2$, the conditional estimates imply an average lower bound for ψ_B of 0.57.

To further validate our findings, we extend the analysis to an international setting covering 20 stock market indices. The results, reported in Table 7, use OECD recession dates and compute the realized conditional Sharpe ratio, the realized price per unit of variance risk, and the expected price per unit of variance risk, as described in Section 3. Consistent with Panel A of Table 5, we find that the Sharpe ratio is higher during recessions for a perfect foresight investor. Similarly, the estimates of the price of variance risk corroborate the evidence in Panel B of the same table. Panel B of Table 7 reports the implied lower bound on the excess volatility ratio based on point estimates of risk prices. We find that this lower bound equals 0.46 and 0.56 during recessions, respectively, closely aligning with our earlier results.

Taken together, the evidence in Tables 5, 6, and 7 delivers a clear message. Even under assumptions that favor the perfect foresight investor entering during recessions, such an investor attains higher Sharpe ratios than an investor in normal times, yet does not earn a higher price per unit of variance risk. Given the weak procyclicality of γ_t, these findings imply that excess volatility must rise sharply in recessions. Imposing zero excess volatility in expansions, we estimate that at least half of the observed market variance in recessions is unpriced.

4.2. Additional macroeconomic evidence

In this section, we further study the cyclical variation in risk prices by linking their fluctuations to variables that capture the state of the financial sector and the overall macro economy. As in the previous section, we first focus on our US sample and thereafter extend the results to a broader international setting.

We start by considering the relation between risk prices and Chicago Fed financial and macroeconomic indicators. The NFCI is the National Financial Conditions Index, which is comprised of several subcategories built to capture risk, credit conditions, and financial and non-financial leverage. High values of the variables are historically associated with tighter-than-average conditions in financial markets, i.e., bad times. The first five rows of Table 8 report the results of regressions on the form:

$$\text{Price of risk}_t = \alpha + \beta \times \text{Financial Risk Indicator}_t + \epsilon_t \tag{17}$$

For the Sharpe ratio results, shown in the first three columns of Table 8, we find that the realized and expected measures are generally negatively related to the financial indicators (Panel A) while the option measure is positively related to most indicators but the positive coefficients are not statistically significant. The last three columns of the table report the results for the price per unit of variance risk. Here, we find that all measures are generally negatively related to the financial risk indicators and ten of the fifteen specifications are statistically significant. These results suggest that the price per unit of variance risk moves strongly procyclically with these risk indicators while the Sharpe ratio moves only mildly with the risk indicators.

To extend the recession results from the previous tables, Panel B of Table 8, we also report results of the relationship between risk prices and the recessions probability of Chauvet and Piger (2008). We find that, when the probability of a recession is high then the price per unit of variance risk is low, lending further evidence of its procyclicality. Results are generally weaker for the Sharpe ratio.

Panel B also reports the results of the Chicago Fed National Activity Index, CFNAI, which is an index built to capture overall economic activity and inflationary pressure. A high value of the CFNAI is generally associated with good economic conditions with high consumption growth, low unemployment, and high industrial production. We find that price per unit of variance risk is positively related to CFNAI and the effect is statistically significant in two of the three measures of the price per unit of variance risk. For the Sharpe ratio, we find no clear relationship between any measure and the CFNAI.

Lastly, Panel B reports results for US consumption growth. Since we have monthly data in the US and quarterly data in our international sample, we divide our analysis into two parts, a US part and an international part. We obtain data on US consumption from the St. Louis Fed database on monthly personal consumption expenditures of non-durable and service goods. We deflate consumption with the CPI. The last column of Table 8 reports the results when regressing the risk prices onto the future eight month consumption growth. Both for the Sharpe ratio and the price per unit of variance risk, we find only mild positive relationships between the measures and consumption growth for the realized and expected measures while the option measures have negative but insignificant coefficients.

In Table 9, we move to an international setting where in Panel A, we report the results of a pooled panel regression of the risk prices onto the subsequent eight quarter consumption growth ($s = 4$ or $s = 8$):

$$\text{Risk price}^i = \alpha_i + \beta \times \text{consumption growth}^i_{q+1,q+s} + \epsilon^i_{q+1,q+s}, \tag{18}$$

where i represents the different countries and Risk price^i is the average of the monthly measures within quarter q. We cluster standard errors by country and quarter and include country fixed effects. We use Final Consumption Expenditure, Real, Unadjusted, Domestic Currency from the IMF database as our proxy for aggregate consumption. The first row reports the results when we pool the raw data and the second row reports results where we standardize both risk prices and consumption growth within each country before pooling the data. In all panel regression specifications, we find that the price per unit of variance risk is positively related to future consumption growth and the slope coefficients are statistically significant for the realized measure. For the Sharpe ratio, we find similar results.

We next study how the risk prices vary with valuation ratios, which are standard measures of the state of the economy in previous asset pricing literature (see e.g. Campbell and Cochrane (1999) and Gormsen and Jensen (2026)). We measure valuation ratios through country-level dividend-price ratios and book-to-market ratios. In Panel B of Table 9, we report the results of panel regressions on the form:

$$\text{Risk price}^i = \alpha_i + \beta \times \text{valuation ratio}^i_t + \epsilon^i_t \tag{19}$$

where we regress the risk prices onto the contemporaneous valuation ratio. We include country fixed effects and cluster standard errors by country and time. The first row for the dividend-price section reports the panel regression results when we pool the raw data for all countries. The second row reports the results where we standardize (mean zero and variance one) both the risk prices and the valuation ratio within each country before pooling. We find that the price per unit of variance risk is negatively related to the contemporaneous level of the dividend-price ratio, suggesting that investors can earn higher price per unit of variance risk in good times when market prices are high. For the Sharpe ratio, we find similar results, but they are mildly weaker for the expected measures.

Finally, we also consider how the risk prices vary with the growth in industrial production. The third part of Table 9 reports panel regressions on the form:

$$\text{Risk price}^i = \alpha_i + \beta \times \text{Industrial production}^i_{t+1,t+s} + \epsilon^i_t \tag{20}$$

where $s = 8$ are months. The first row of the industrial production part of the table reports results using the raw data and the second row reports the results when we standardize the data input within country before pooling the data. Our data on industrial production is from the OECD database. We find that the price per unit of variance risk is positively related to future growth in industrial production in all our regression specifications and the slopes are all statistically significant. The coefficients are also positive for the Sharpe ratio but less statistically significant for the expected measure. In these panel regressions, we again add country fixed effects and cluster standard errors by country and time. These results show that the risk prices are high in good times when economic activity as measured by the growth in industrial production is high.

4.3. Conditional excess volatility and real investments

Table 12 examines whether conditional excess volatility predicts subsequent real investment. The estimates show that investment growth declines following periods in which the price per unit of variance risk is unusually low, corresponding to episodes of high conditional excess volatility. The effect is strongest over the subsequent one to two quarters and remains economically negative at longer horizons. These results are robust to controlling for standard proxies for uncertainty and macroeconomic conditions, including the VIX, recession probabilities, financial conditions, valuation ratios, and credit spreads. This evidence suggests that conditional excess volatility contains information about firms' investment decisions that is not captured by existing measures of aggregate risk or the business cycle.

These findings connect to a broad literature studying the effects of uncertainty on investment. A central insight of this literature is that heightened uncertainty increases the value of waiting, leading firms to postpone irreversible investment (Bernanke, 1983; Bloom, 2009). Our results point to a complementary mechanism. While previous work largely treats uncertainty as synonymous with aggregate volatility or macroeconomic risk, the evidence here indicates that the component of volatility that is orthogonal to the stochastic discount factor also has real effects. Because conditional excess volatility does not command a risk premium by construction, the decline in investment cannot be attributed solely to changes in priced risk. Instead, firms appear to respond to fluctuations in uncertainty that are not reflected in expected returns. This interpretation complements recent evidence that volatility-managed strategies earn abnormal returns because periods of high volatility are disproportionately driven by uncompensated risk (Moreira and Muir, 2017), while extending the implications of excess volatility beyond financial markets to firms' real investment decisions. More broadly, the results suggest that distinguishing between priced and unpriced sources of aggregate risk is important not only for understanding asset prices but also for explaining fluctuations in capital accumulation.

5. Consumption-Based Evidence

Consumption-based models offer an alternative way to investigate cyclical patterns in risk pricing and excess volatility. As discussed in Section 2, γ_t has a natural interpretation as the beta of the SDF projected onto the market. Therefore, we now write the price per unit of variance risk as

$$\gamma_t = -\text{Cov}_t(\log M_{t+1}, \log R_{m,t+1})/\sigma^2_{m,t}, \tag{21}$$

where $\sigma^2_{m,t}$ is the variance of log market returns. Applying this identity to the most widely used consumption-based stochastic discount factors, we get:

	M_{t+1}	γ_t
Power utility:	$\delta_c\,(\Delta C_{t+1})^{-\gamma_c}$	$\gamma_c \rho_{m,c,t} \frac{\sigma_{c,t}}{\sigma_{m,t}}$
Recursive utility:	$\delta_c^{\theta_c} R_{w,t+1}^{\theta_c - 1} (\Delta C_{t+1})^{-\frac{\theta_c}{\phi_c}}$	$\frac{\theta_c}{\phi_c} \rho_{m,c,t} \frac{\sigma_{c,t}}{\sigma_{m,t}} + (1-\theta_c)\rho_{m,w,t} \frac{\sigma_{w,t}}{\sigma_{m,t}}$

where δs are time-preference parameters, γ_c is the relative risk aversion over consumption growth, ΔC_{t+1}, and $\theta_c = (1-\gamma_c)/(1-\frac{1}{\phi_c})$ where ϕ_c is the intertemporal elasticity of substitution. $R_{w,t+1}$ is the return on the wealth portfolio. ρs are correlations between the market and either the consumption growth or wealth portfolio, and σs are their respective volatilities.

Table 10 reports empirical quantities of the correlations and volatilities that enter in the utility specifications. Columns 1 and 2 report the empirical correlations. There is little difference in the correlation between the log consumption growth and the log stock market returns during recessions and normal times, 0.1428 and 0.1771 respectively. The result that the empirical correlation is slightly higher in normal times is consistent with the results in Duffee (2005), who argues that the ”composition effect” drives the correlation. When stock market wealth composes a high fraction of total wealth then consumption and stock market returns are more tightly intertwined which makes them more highly correlated. Xu (2020) also finds supporting evidence of a procyclical behavior of the correlation between stock market returns and consumption growth. She shows that the procyclicality comes from a procyclical correlation between dividend growth and consumption growth. Columns 3 through 5 report the volatilities. Consumption volatility is 30.26% higher in recessions than in normal times, which is much lower than for the stock market volatility that is 69.84% higher in recessions.

Plugging the non-recession numbers in Table 10 into the power utility specification of the price per unit of variance risk, we get

$$\gamma_t = \underbrace{100}_{\gamma_c} \times 0.1771 \times \frac{0.0042}{0.0663} = 1.1228, \tag{22}$$

where $\gamma_c = 100$ to get reasonable size risk prices and average expected excess returns, see e.g. Mehra and Prescott (1985). Doing a similar exercise using NBER recession numbers, we find that the price per unit of variance risk in recessions is 0.6948. These results suggest that the price per unit of variance risk is procyclical, it is almost twice as high in normal times as opposed to in recessions. We notice that, while the price per unit of variance risk is higher in normal times, the risk premium is still highest in recessions. We find the risk premium as $\gamma_t \sigma^2_{m,t}$, which is 0.0352 for recession periods and 0.0197 in normal times. These numbers correspond to about 14% annualized risk premium in recessions and 8% in normal times.

Turning to recursive utility, we follow Bansal and Yaron (2004) and set $\gamma_c = 10$ and allow for the intertemporal elasticity of substitution to be greater than one, $\phi_c = \frac{1}{0.75}$. Plugging the numbers in Table 10 into the recursive utility specification of the price per unit of variance risk, we find that it is 6.79 in normal times and 5.40 during recessions. These results further suggest that the price per unit of variance risk is indeed procyclical. We have tried a wide grid of combinations of the preference parameters for γ_c and ϕ_c and find that price per unit of variance risk is procyclical, except for a few cases with very low relative risk aversion.

Our findings in this section further corroborate our empirical result that the price per unit of variance risk is procyclical. It further highlights excess volatility as a potential mechanism to explain market variation. We explore a model of excess volatility theoretically in the next section. That is, as Table 10 shows, what happens when going from normal times to recession periods is that the stock market volatility tends to increase much more than the volatility of fundamentals like consumption and the return on the wealth portfolio. There is a disconnect between stock markets and fundamentals in recessions and this disconnect is not offset by higher correlations, on the contrary, correlations seem to be constant over time.

6. A Model of Conditional Excess Volatility

Next, we reconcile our empirical findings on the cyclicality of the price per unit of variance risk with stylized facts about the standard equity term structure. We consider a model building off of that of Lettau and Wachter (2007), but with a new assumption that there are "non-fundamental" shocks to dividend claims. These non-fundamental shocks increase in importance in bad times, when the fundamental price of risk is high. We will see that this model will produce a countercyclical Sharpe ratio, as is standard. But it will also feature a *lower* price per unit of variance risk in bad times, and thus a procyclical price of this risk. Finally, the greater exposure of long-horizon claims to unpriced discount-rate shocks will produce a downward-sloping term structure of Sharpe ratios, consistent with the Lettau and Wachter (2007) model and with empirical estimates (van Binsbergen and Koijen, 2017; Cejnek and Randl, 2020; Golez and Jackwerth, 2024).[6]

6.1. Model setup

The aggregate dividend is denoted by D_t, and let $d_t = \log D_t$. We assume that log dividend growth follows

$$\Delta d_{t+1} = g - \frac{1}{2}x_t^2\sigma_z^2 + \sigma_d\varepsilon_{d,t+1} + x_t\sigma_z\varepsilon_{z,t+1}, \tag{23}$$

where $\varepsilon_{d,t+1}$ and $\varepsilon_{z,t+1}$ are standard normal and independent of each other and over time. Relative to the specification in Lettau and Wachter (2007), we include an additional shock $x_t\sigma_z\varepsilon_{z,t+1}$, whose volatility is time-varying and increasing in x_t. This x_t variable will also, for parsimony, represent the price of risk.[7] The shock $\varepsilon_{d,t+1}$ will be priced (i.e., it will enter the SDF), while the shock with time-varying volatility will not be. This should be thought of as a stripped-down way to model the idea that returns on dividend strips (and the market) include additional "non-fundamental" volatility in bad times, reverse-engineered here by including an unpriced shock in dividend growth that becomes more important in bad times.

As above, the price of risk is driven by a single state variable x_t, which follows

$$x_{t+1} = (1-\phi_x)\overline{x} + \phi_x x_t + \sigma_x\varepsilon_{x,t+1}, \tag{24}$$

where $\varepsilon_{x,t+1}$ is i.i.d. standard normal and independent of $\varepsilon_{d,t+1}$ and $\varepsilon_{z,t+1}$ and where $\overline{x} > 0$, $\phi_x \in (0,1)$.

[6] In the main version of the model, we assume that shocks to the fundamental price of risk are unpriced, for analytical convenience. But if, in addition, we included a small price of risk on discount-rate shocks, one could reconcile this model with the countercyclical term premium as in Gormsen (2021). We plan to do so in a full quantitative version of the model.

[7] As an additional twist relative to Lettau and Wachter (2007), we also rule out time variation in the conditional mean of (exponentiated) dividend growth. Including such time variation would allow for a downward-sloping term structure of expected returns (rather than a constant term structure of expected returns but downward-sloping Sharpe ratios and CAPM alphas, as ours will feature). Since this is relatively unimportant for our analysis, we simplify by omitting such time variation.

The log stochastic discount factor $m_{t+1} = \log M_{t+1}$ is directly specified as

$$m_{t+1} = -r^f - \frac{1}{2}x_t^2 - x_t \varepsilon_{d,t+1}, \tag{25}$$

where r^f is the constant log risk-free rate. Intuitively, investors dislike exposure to "fundamental" dividend-growth shocks $\varepsilon_{d,t+1}$, and the degree to which they dislike this exposure is governed by risk aversion (the conditional price of risk) x_t. All other risks are unpriced directly. All the main results would carry through in more realistic cases in which fundamental discount-rate shocks (i.e., shocks to x_t) are priced (with a smaller price on this risk than on fundamental risk); we omit this for now for analytical clarity.

6.2. Solution and implications

We now solve explicitly for the prices and returns of zero-coupon equity (i.e., n-maturity dividend claims).[8] The price of the n-maturity claim at time t is $P_{n,t}$, and let $p_{n,t} = \log P_{n,t}$. One-period returns are $R_{n,t+1} = P_{n-1,t+1}/P_{n,t}$. Since $\mathbb{E}_t[M_{t+1}R_{n,t+1}] = 1$, we have the following recursive relation for prices:

$$P_{n,t} = \mathbb{E}_t[M_{t+1}P_{n-1,t+1}], \tag{26}$$

with boundary condition $P_{0,t} = D_t$ given that the dividend is paid out at maturity.

Guess a log-linear solution for the price-dividend ratio:

$$\frac{P_{n,t}}{D_t} = \exp(A_n + B_{x,n}x_t)\,. \tag{27}$$

Under this conjecture, the price-dividend ratio is

$$\frac{P_{n,t}}{D_t} = \mathbb{E}_t\left[M_{t+1}\frac{D_{t+1}}{D_t}\exp(A_{n-1} + B_{x,n-1}x_{t+1})\right]. \tag{28}$$

Using the assumed conditional log-normality in (23)–(24), we can match coefficients of (27)

[8] The price and return for aggregate equity then follows straightforwardly from the zero-coupon solutions, but in order to examine intuition, we maintain focus on the zero-coupon claims.

and (28) to obtain

$$A_n = A_{n-1} - r^f + g + B_{x,n-1}(1-\phi_x)\overline{x} + \frac{1}{2}\sigma_d^2 + \frac{1}{2}B_{x,n-1}^2\sigma_x^2, \tag{29}$$

$$B_{x,n} = B_{x,n-1}\phi_x - \sigma_d = -\frac{1-\phi_x^n}{1-\phi_x}\sigma_d, \tag{30}$$

with boundaries $A_0 = B_{x,0} = 0$. This verifies the conjecture. Note that $B_{x,n} < 0$ for all n, so that the price-dividend ratio decreases (times are bad) when the price of risk increases.

The log return on the strip of maturity n is then

$$\begin{aligned} r_{n,t+1} &= \log\left(\frac{P_{n-1,t+1}}{D_{t+1}}\frac{D_t}{P_{n,t}}\frac{D_{t+1}}{D_t}\right) \\ &= g - \frac{1}{2}x_t^2\sigma_z^2 + \sigma_d\varepsilon_{d,t+1} + x_t\sigma_z\varepsilon_{z,t+1} + A_{n-1} + B_{x,n-1}x_{t+1} - A_n - B_{x,n}x_t. \end{aligned} \tag{31}$$

The conditional variance of this log return follows as

$$\sigma_{n,t}^2 = \text{Var}_t(r_{n,t+1}) = \sigma_d^2 + \sigma_z^2 x_t^2 + |B_{x,n-1}|^2\sigma_x^2. \tag{32}$$

The excess expected return, meanwhile, is

$$\mathbb{E}_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_{n,t}^2 = -\text{Cov}_t(r_{n,t+1}, m_{t+1}) = \sigma_d x_t, \tag{33}$$

so the term structure of expected returns is flat.

Putting (32) and (33) together, the Sharpe ratio is

$$\text{SR}_{n,t} \equiv \frac{\mathbb{E}_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_{n,t}^2}{\sigma_{n,t}} = \frac{\sigma_d x_t}{\sqrt{\sigma_d^2 + \sigma_z^2 x_t^2 + B_{x,n-1}^2\sigma_x^2}}. \tag{34}$$

Note first that this is decreasing in maturity n, so we obtain a downward-sloping term structure of Sharpe ratios. This holds because longer-maturity claims are more exposed to discount-rate risk. In addition, the Sharpe ratio is countercyclical:

$$\frac{\partial \text{SR}_{n,t}}{\partial x_t} \propto \left(\sigma_d^2 + B_{x,n-1}^2\sigma_x^2\right) > 0. \tag{35}$$

The countercyclical price of risk passes through to generate a countercyclical Sharpe ratio, as is standard.

But the ratio of expected returns to *variance*, meanwhile, is

$$\gamma_{n,t} \equiv \frac{\mathbb{E}_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_{n,t}^2}{\sigma_{n,t}^2} = \frac{\sigma_d x_t}{\sigma_d^2 + \sigma_z^2 x_t^2 + |B_{x,n-1}|^2 \sigma_x^2}, \tag{36}$$

which varies with x_t according to

$$\frac{\partial \gamma_{n,t}}{\partial x_t} \propto \left(\sigma_d^2 - \sigma_z^2 x_t^2 + B_{x,n-1}^2 \sigma_x^2\right). \tag{37}$$

This value can be either positive or negative, and it will be negative if and only if

$$\sigma_z |x_t| > \sigma_d \sqrt{1 + \sigma_x^2 \left(\frac{1 - \phi_x^{n-1}}{1 - \phi_x}\right)^2}. \tag{38}$$

So for x_t large enough, we obtain a *procyclical* price per unit of variance risk: further positive shocks to x_t increase non-fundamental return volatility enough to offset the increase in expected returns. We will maintain the assumption that the steady-state value $\overline{x}$ is large enough to ensure such procyclicality around steady state; we show later that this assumption is consistent with reasonable calibrations of the model. And in part of the state space, there will be a countercyclical $\gamma_{n,t}$ (in good times, when x_t is small), as we observe in the time-series data. But in bad enough times, non-fundamental return volatility is large enough to make the price per unit of variance risk decrease given further increases in x_t.

Intuitively, expected returns go up with the price of risk, but the importance of non-fundamental risk for returns also increases. While return volatility increases, the beta of the SDF onto the market decreases during these times because of the rise in non-fundamental return risk. In other words, return volatility has three components: fundamental dividend volatility (which is constant, σ_d), fundamental discount-rate volatility (also constant, σ_x), and non-fundamental volatility (which increases in x_t). An increase in x_t therefore increases "pure" market risk that is not fully connected to fundamentals, increasing market variance without passing through one-for-one to expected returns. While this is sufficient to generate a procyclical $\gamma_{n,t}$ that decreases in x_t (at least for large x_t), this non-fundamental volatility effect is not strong enough to obtain a procyclical Sharpe ratio. Further, the greater exposure of long-maturity claims to discount-rate risk means that their volatility increases without changing their expected return, generating a downward-sloping Sharpe ratio of dividend claims.

As a result, this stylized model shows how our findings about the cyclicality of $\gamma_{n,t}$ connect to facts about the dividend term structure. We obtain *both* (i) a lower beta of the SDF onto the market in bad times (our main stylized fact), and (ii) a downward-sloping Sharpe ratio of dividend claims by maturity, both through the same channel (non-fundamental return risk). Meanwhile, our setting maintains the usual countercyclical Sharpe ratio.

With respect to the *cyclicality* of the term structure, speaking to the countercyclicality documented by Gormsen (2021) could be achieved by assuming, similar to Gormsen, that the discount-rate shock $\varepsilon_{x,t+1}$ also enters into the SDF, with a small average price of risk for this shock but an increase in the quantity of this risk in bad times.

6.3. Illustrative quantification and simulations

The previous subsection shows the model's qualitative consistency with the paper's main results. We now examine whether the model as specified is *quantitatively* consistent with the empirical estimates, using a set of calibrated model simulations. We calibrate the model using prior literature's estimates when possible. For the new parameters in our model related to non-fundamental risk, we choose parameters to generate a reasonable standard deviation and Sharpe ratio of long-horizon dividend claims.[9]

Parameter values are shown in Panel A of Table 11. We choose ϕ_x, g, r^f, σ_d, and σ_x following Gormsen (2021), and parameters are also close in value to those in Lettau and Wachter (2007). We choose $\overline{x}$ to generate a constant 7% risk premium on each maturity's dividend claim (and the market as a whole), which is close to the value in Lettau and Wachter (2007).[10] Finally, we select σ_z to generate a long-horizon dividend-strip Sharpe

[9]In future iterations, we plan to conduct formal statistical estimation of model parameters via maximum likelihood.

[10]We choose a very slightly higher value so that the condition in (38) is satisfied in steady state, which requires $\overline{x} \geq 0.65$ given other parameters, though this is relatively unimportant.

ratio of 0.35.

Simulation results are presented in Panel B of Table 11. The first sub-panel shows that unconditional average returns, standard deviations, and Sharpe ratios are reasonable for all maturities' claims, and the price per unit of variance risk $\gamma_{n,t}$ is on average quantitatively close to the average expected γ_t and the average option-implied γ_t from previous sections. Sharpe ratios are also downward-sloping by maturity, consistent with the data and the analytical results in the previous subsection.

The next two sub-panels show results for the cyclicality of the Sharpe ratio and the price per unit of variance risk. "Very bad times" can be thought of as similar to the recessions in our sample, which account for roughly 10% of the data, but we also show results split by above and below the steady-state value of the state variable x_t. As in the data, the second sub-panel shows that the Sharpe ratio is countercyclical: it increases for each maturity in bad vs. good times, and particularly so in deep "recessions." Note as well that the term structure of the Sharpe ratios is countercyclical: in bad times, long-maturity Sharpe ratios are close to short-maturity Sharpe ratios, while the downward slope in maturity is steeper in good times. This mirrors the finding in Gormsen (2021) that the slope of the term structure of risk premia is countercyclical. Our simpler setting — with constant expected dividend growth and no price associated with price-of-risk shocks — does not generate any term structure or cyclicality in risk premia themselves, but the patterns for risk-adjusted returns mirror Gormsen's results at a conceptual level.[11]

The price per unit of variance risk, meanwhile, is procyclical: as shown in the last sub-panel, $\gamma_{n,t}$ is meaningfully lower in bad times than in good times (particularly in very bad times). And the spread between very bad times and good times is quantitatively similar to the spread observed in the data for conditional expected γ_t and the conditional option-implied γ_t. As a result, we conclude that *all* the key features documented empirically are matched reasonably well quantitatively by our model, which also produces a term structure of risk-adjusted returns consistent with known facts about the dividend claims term structure.

[11]Including these additional time-varying features would be straightforward and would not change the model's conclusions.

7. Term structure evidence from option portfolios

The preceding model shows that a procyclical γ_t implies a downward-sloping term structure of Sharpe ratios on dividend claims. In this section, we provide complementary evidence from option markets that serves as an out-of-sample test of the procyclicality result. The key insight is that option portfolios designed to hold fixed the *quantity* of risk have expected returns that depend only on γ_t. If γ_t covaries negatively with the SDF — that is, if it is higher in good times — then multiperiod versions of these portfolios provide a hedge against bad states, and their expected returns should accordingly decrease with the horizon.

Setup. Consider a conditionally log-normal economy in which the log SDF loads on the market return with coefficient γ_t, and there is an additional orthogonal SDF component η_{t+1} (see Section B for formal details). An Arrow–Debreu security for return state ω has a log expected excess return of $\gamma_t(\omega - \mu_{R,t} + \frac{1}{2}\sigma^2_{\varepsilon,t})$. Now consider a binary bet that goes short one unit of the Arrow–Debreu security for return state ω_1 and long one unit for state $\omega_2 > \omega_1$. This strategy has a fixed payoff of ± 1, and its log expected return is

$$\log \mathbb{E}_t[R_{\omega,t+1}] = \gamma_t(\omega_2 - \omega_1). \tag{39}$$

Because the strategy holds fixed the return outcomes ω_1 and ω_2, its expected return depends only on the *price* per unit of variance risk γ_t, not on the quantity of risk σ_t^2. This isolation of the pricing component is what makes option portfolios useful for our purposes.

Term structure prediction. To see why equation (39) implies a term structure prediction, step back one period to $t-1$ and consider the two-period return on the same strategy. The expected return from t to the option's expiration $t+1$ will be higher following a positive shock to γ_t. If $\mathrm{Cov}_{t-1}(\gamma_t, M_t) < 0$ — that is, if γ_t is higher in good times — then the strategy's return from $t-1$ to t is positively correlated with bad states: the strategy provides a hedge. As a result, the two-period expected return should be lower than the one-period expected return, implying a downward-sloping term structure of risk prices by horizon.

Conversely, if γ_t were countercyclical (higher in bad times), the strategy would load positively on bad states and command a *higher* expected return at longer horizons. The sign of the term structure slope therefore provides a direct test of the cyclicality of γ_t.

Empirical implementation. We use S&P 500 index options from the OptionMetrics database for the period 1996–2018. For each observed expiration date T and initial trading date 0, we define a set of equally spaced return states based on 2-percentage-point log excess return bins within $\pm 10\%$ of the risk-free rate. We estimate the risk price $\phi_\kappa = \mathbb{E}[M_T | R_{m,T} = \omega_1] / \mathbb{E}[M_T | R_{m,T} = \omega_2]$ at each horizon $\kappa = T - t$ (in weeks) by constructing risk-neutral probabilities from the option-implied distribution and using a GMM procedure that accounts for option price measurement error through lagged instruments. Full details of the estimation are in Section B.2.

Results. Figure 3 presents the estimated risk prices by horizon. The term structure is significantly downward-sloping: the estimated ϕ_κ declines from roughly 1.30 at the one-week horizon to approximately 1.05 at twelve weeks. This pattern implies that a lower price of risk is needed to rationalize the returns on fixed-quantity-of-risk bets at longer horizons, or equivalently, that these bets earn lower expected returns at longer horizons.

Following the logic above, the downward slope is consistent with $\text{Cov}_{t-1}(\gamma_t, M_t) < 0$: the price per unit of risk is higher in good times, so that the binary bet strategy provides a hedge against bad states at longer horizons. This is precisely what the procyclicality of γ_t — established in Section 4 using entirely different data and methodology — would predict. The option evidence thus serves as an independent, out-of-sample validation of our main result.

The evidence from this estimation is admittedly suggestive rather than definitive, given the relatively short horizons (1–12 weeks) available in our data. Whether this weekly-to-monthly frequency evidence extends to business-cycle frequencies remains an open question. Nonetheless, the qualitative consistency with the model's predictions and the main empirical results is encouraging.

8. Conclusion

We propose a simple, model-free benchmark for identifying excess volatility in equity markets. By decomposing market return variance into a component spanned by the stochastic discount factor (SDF) and an orthogonal residual, we define "conditional excess volatility" as the share of return variance that cannot command a risk premium. The key statistic $\gamma_t = \mu_t/\sigma_t^2$ identifies the fraction of variance that is priced through the decomposition $\gamma_t = \widetilde{\gamma}_t(1 - \psi_t)$, where $\widetilde{\gamma}_t$ is the price of spanned SDF risk and ψ_t is the fraction of total market variance that is orthogonal to the SDF. Estimating γ_t using realized returns, predictive models of expected returns, and option-implied measures across 20 equity markets, we find that it does not rise in recessions and often declines. Because standard asset-pricing models predict that the price of spanned SDF risk is countercyclical, this evidence implies that the share of variance orthogonal to the SDF increases sharply in downturns. In quantitative terms, at least half of total market variance during recessions appears to reflect excess volatility rather than compensation for SDF risk.

These findings provide a unified interpretation of several empirical puzzles. When excess volatility rises in bad times, increases in market variance need not translate into proportionally higher expected returns, rationalizing why optimal portfolio weights decline during volatility spikes despite elevated Sharpe ratios, why volatility-managed strategies earn alpha, and why the term structure of Sharpe ratios on equity claims slopes downward. A simple asset-pricing model with non-fundamental dividend shocks reproduces these patterns and links them to the broader equity term structure. More broadly, our results suggest that fluctuations in market volatility partly reflect variation unrelated to marginal utility, especially during crises.

References

AIT-SAHALIA, Y. AND A. W. LO (2000): "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 94, 9—-51.

BANSAL, R. AND A. YARON (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, 59, 1481–1509.

BATES, D. S. (2000): "Post 87 Crash Fears in the S&P 500 Futures Options Market," *Journal of Econometrics*, 94, 181–238.

BERKOWITZ, J. (2001): "Testing Density Forecasts, With Applications to Risk Management," *Journal of Business and Economic Statistics*, 19, 465–474.

BERNANKE, B. S. (1983): "Irreversibility, Uncertainty, and Cyclical Investment," *The Quarterly Journal of Economics*, 98, 85–106.

VAN BINSBERGEN, J. H. AND R. S. J. KOIJEN (2017): "The Term Structure of Returns: Facts and Theory," *Journal of Financial Economics*, 124, 1–21.

BLISS, R. R. AND N. PANIGIRTZOGLOU (2004): "Option Implied Risk Aversion Estimates," *Journal of Finance*, 59, 407–446.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

BREEDEN, D. T. AND R. H. LITZENBERGER (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *The Journal of Business*, 51, 621—-651.

CAMPBELL, J. Y. AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

CAMPBELL, J. Y., S. GIGLIO, R. TURLEY, AND C. POLK (2018): "An intertemporal CAPM with stochastic volatility," *Journal of Financial Economics*, 128, 207–233.

Campbell, J. Y. and R. J. Shiller (1987): "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95, 1062–1088.

Cejnek, G. and O. Randl (2020): "Dividend Risk Premia," *Journal of Financial and Quantitative Analysis*, 55, 1199–1242.

Chauvet, M. and J. Piger (2008): "A comparison of the real-time performance of business cycle dating methods," *Journal of Business and Economic Statistics*, 26, 42–49.

Duffee, G. R. (2005): "Time Variation in the Covariance between Stock Returns and Consumption Growth," *Journal of Finance*, 60, 1673—1712.

Fama, E. F. (1968): "Risk, Return and Equilibrium: Some Clarifying Comments," *Journal of Finance*, 23, 29–40.

Figlewski, S. (2018): "Risk Neutral Densities: A Review," *working paper*.

Friend, I. and M. E. Blume (1975): "The Demand for Risky Assets," *American Economic Review*, 65, 900–922.

Golez, B. and J. Jackwerth (2024): "Holding Period Effects in Dividend Strip Returns," *Review of Financial Studies*, 37, 3188–3215.

Gormsen, N. J. (2021): "Time Variation of the Equity Term Structure," *Journal of Finance*, 76, 1959–1999.

Gormsen, N. J. and C. S. Jensen (2024): "Conditional Risk," *Journal of Financial Economics*, 162.

——— (2026): "Higher Moment Risk," *Journal of Finance, forthcoming*.

Jackwerth, J. (2000): "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.

Jackwerth, J. C. (2004): "Option-Implied Risk-Neutral Distributions and Risk Aversion," *CFA Institute Research Foundation of AIMR Publications*, 1–86.

JENSEN, C. S. (2018): "Essays on Asset Pricing," *PhD dissertation, Copenhagen Business School*.

JENSEN, C. S., D. LANDO, AND L. H. PEDERSEN (2019): "Generalized Recovery," *Journal of Financial Economics*, 133(1), 154–174.

KELLY, B. AND S. PRUITT (2013): "Market Expectations in the Cross-Section of Present Values," *Journal of Finance*, 68, 1721–1756.

LAZARUS, E. (2022): "Horizon-Dependent Risk Pricing: Evidence from Short-Dated Options," *Working paper*.

LEROY, S. F. AND R. D. PORTER (1981): "The Present Value Relation: Tests Based on Implied Variance Bounds," *Econometrica*, 49, 555–574.

LETTAU, M. AND S. LUDVIGSON (2001): "Consumption, Aggregate Wealth, and Expected Stock Returns," *The Journal of Finance*, 56, 815—-849.

LETTAU, M. AND J. A. WACHTER (2007): "Why is long-horizon equity less risky? a duration-based explanation of the value premium," *The Journal of Finance*, 62, 55–92.

LUSTIG, H. AND A. VERDELHAN (2012): "Business cycle variation in the risk-return trade-off," *Journal of Monetary Economics*, 59, S35–S49.

MEHRA, R. AND E. C. PRESCOTT (1985): "The equity premium: A puzzle," *Journal of monetary Economics*, 15, 145–161.

MOREIRA, A. AND T. MUIR (2017): "Volatility-Managed Portfolios," *Journal of Finance*, 72, 1611—-1644.

ROSENBERG, J. V. AND R. F. ENGLE (2002): "Empirical Pricing Kernels," *Journal of Financial Economics*, 64, 341—-372.

SCHREINDORFER, D. AND T. SICHERT (2023): "Volatility and the Pricing Kernel," *working paper*.

SHILLER, R. J. (1981): "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes In Dividends?" *American Economic Review*, 71, 421–436.

XU, N. R. (2020): "Procyclicality of the Comovement between Dividend Growth and Consumption Growth," *Journal of Financial Economics, forthcoming*.

Figure 2. Estimates of Risk Prices. Panel A shows the Sharpe ratio of the S&P 500 index. Panel B shows the price per unit of variance risk, γ_t. We standardize the measures to make them comparable in the figure. Risk prices are computed using the methods described in Section 3. Shaded area is NBER recession periods.

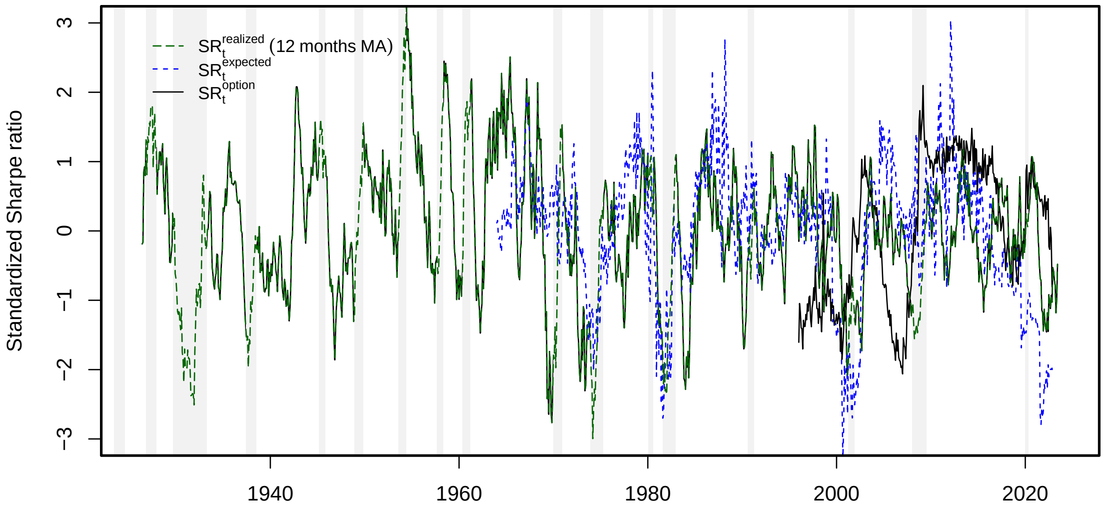

Panel A. Sharpe ratio

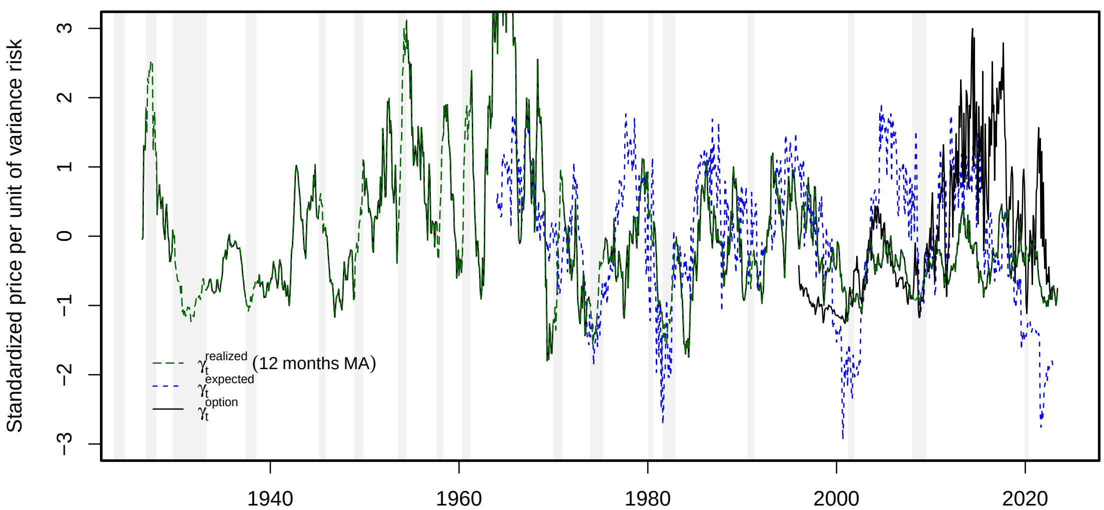

Panel B. Price per unit of variance risk

Figure 3. Estimates of Risk Prices by Horizon. This figure shows the option-implied price of risk by horizon. Point estimates are constructed using two-step GMM, using the five-day-lagged observation as an instrument as described in Appendix B.2, on the sample counterparts of the moment conditions described in the appendix in order to minimize forecast error. The price of risk parameter is constrained to be equal for all days within a given weekly horizon to expiration. Error bars show 95% confidence intervals, constructed using the procedure described in Appendix B.2.

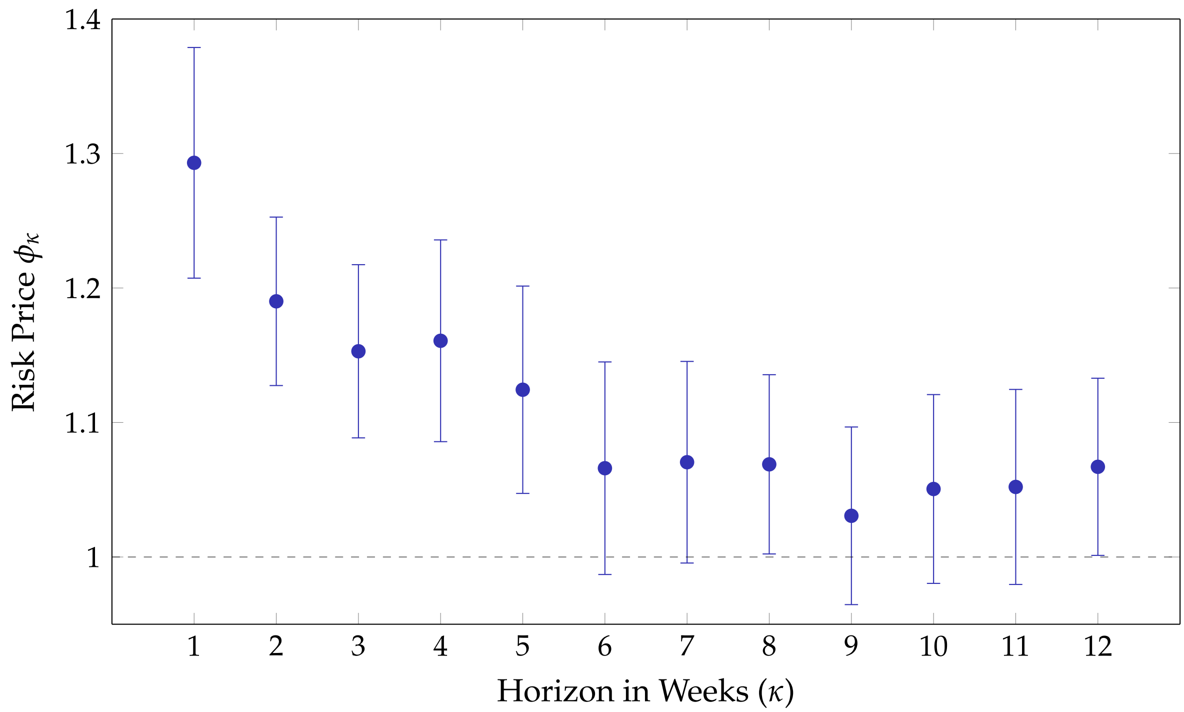

Table 4
Pairwise correlations of risk prices

This table reports the pairwise correlations of the price per unit of variance risk (γ_t) and the Sharpe ratio (SR_t) for the S&P 500 index. We compute four measures for each of the risk prices: (i) a 12-months moving average of the monthly realized measure $\gamma_t^{\text{realized}}$, using daily returns to compute conditional expected excess returns and volatilities, (ii) $\gamma_t^{\text{expected}}$, as an expected measure using Kelly and Pruitt (2013) for expected excess returns and an AR(1) process for the variance, (iii) γ_t^{option}, an option implied measure estimated from options written on the S&P 500 index, and (iv) γ_t^{PCA}, the first principal component of the three previous measures. Details on how we compute risk prices can be found in Section 3.

	$\gamma_t^{\text{expected}}$	γ_t^{option}	γ_t^{PCA}	SR_t^{realized}	SR_t^{expected}	SR_t^{option}	SR_t^{PCA}
$\gamma_t^{\text{realized}}$	0.42	0.29	0.46	0.84	0.30	-0.09	0.28
$\gamma_t^{\text{expected}}$		0.27	0.66	0.43	0.91	-0.03	0.66
γ_t^{option}			0.78	0.28	0.22	0.54	0.51
γ_t^{PCA}				0.42	0.59	0.25	0.77
SR_t^{realized}					0.39	0.08	0.41
SR_t^{expected}						0.13	0.80
SR_t^{option}							0.54

Table 5

Cyclicality in risk prices

This table compares the price of risk for the US stock market in good times (subscript "G") to the months in recessions **after** the stock market reached its low during the recession (subscript "B"). Specifically, we compare the prices of risk for an investor who only invests in non-NBER recession months to that of an investor who has perfect foresight during NBER recessions in the sense that she can pinpoint when the market has reached its low. This investor buys the market at its low and holds the market for twelve months. We compute unconditional measures by bundling monthly returns based on each trading strategy and compute within bundle expected excess returns and variance. The rows report the differences between recession periods and normal times. A positive value means that the point estimate is higher in bad times than in good times. *t*-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors. * indicates statistical significance at the 10% level.

	Post 1928 sample	Post 1964 sample	Post 1996 sample
Panel A: Market Sharpe ratio			
Unconditional	0.34*	0.44*	0.29
s.e.	0.10	0.16	0.27
$SR_B^{realized} - SR_G^{realized}$	1.78*	1.13*	0.43
	4.45	2.19	0.52
$SR_B^{expected} - SR_G^{expected}$		−0.17*	−0.26
		−1.69	−1.15
$SR_B^{implied} - SR_G^{implied}$			0.10*
			1.91
Panel B: Price per unit of market variance risk			
Unconditional	4.07*	7.50	3.44
s.e.	2.44	3.76	5.95
$\gamma_B^{realized} - \gamma_G^{realized}$	14.36*	−1.79	−4.78
	1.75	−0.31	−0.96
$\gamma_B^{expected} - \gamma_G^{expected}$		−1.77*	−2.23*
		−3.65	−2.46
$\gamma_B^{implied} - \gamma_G^{implied}$			−0.65*
			−2.06
Panel C: Sample statistics			
No. recession periods	16	8	3
No. recession months after market low	192	96	36
No. normal times months	865	556	277

Table 6

Required excess volatility

Based on the point estimates of risk prices used for Table 5, this table reports empirical lower bound estimates of

$$\psi_B = 1 - \frac{\gamma_B}{\gamma_G}\frac{1-\psi_G}{k}. \tag{40}$$

where γ_B is the estimate of the price per unit of market variance risk in recessions (bad times) and γ_G is the non-recession estimate. The panels report the ratio of the risk prices, the minimum excess volatility ratio in good times, ψ_G, and the minimum excess volatility ratio in bad times – the fraction of the total market variance in good and bad times that is unrelated to fundamentals. Given the ratio $\frac{\gamma_B}{\gamma_G}$ and a value of k, we find the lowest possible $\psi_G \geq 0$ such that $\psi_B \geq 0$ in Equation (40). At this minimum value of ψ_G, we then compute the minimum ψ_B. The quantity $k = \widetilde{\gamma}_B/\widetilde{\gamma}_G$ denotes the relative price of SDF-relevant risk in bad times over good times. If $k = 2$ then the price of SDF-relevant risk doubles in bad times. We report 90% bootstrapped confidence intervals and * indicates statistical significance at the 10% level.

	Post 1928 sample			Post 1964 sample			Post 1996 sample		
	$k=1$	$k=2$	$k=4$	$k=1$	$k=2$	$k=4$	$k=1$	$k=2$	$k=4$
Panel A: Unconditional									
γ_B/γ_G	2.12	2.12	2.12	3.29	3.29	3.29	1.83	1.83	1.83
Minimum ψ_G	0.53*	0.06	0	0.70*	0.39	0	0.45*	0	0
90% CI, ψ_G	[0.23, 0.76]	[0.00, 0.52]	[0.00, 0.04]	[0.50, 0.85]	[0.00, 0.70]	[0.00, 0.39]	[0.14, 0.69]	[0.00, 0.38]	[0.00, 0.00]
Minimum ψ_B	0	0	0.47	0	0	0.18	0	0.08	0.54*
90% CI, ψ_B	[0.00, 0.00]	[0.00, 0.35]	[0.00, 0.67]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.50]	[0.00, 0.00]	[0.00, 0.42]	[0.19, 0.71]
Panel B: $\gamma_t^{realized}$									
γ_B/γ_G	1.55	1.55	1.55	0.92	0.92	0.92	0.70	0.70	0.70
Minimum ψ_G	0.36*	0	0	0	0	0	0	0	0
90% CI, ψ_G	[0.06, 0.58]	[0.00, 0.16]	[0.00, 0.00]	[0.00, 0.46]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.46]	[0.00, 0.00]	[0.00, 0.00]
Minimum ψ_B	0	0.22	0.61*	0.08	0.53*	0.77*	0.30	0.65*	0.82*
90% CI, ψ_B	[0.00, 0.00]	[0.00, 0.46]	[0.41, 0.73]	[0.00, 0.46]	[0.09, 0.74]	[0.54, 0.87]	[0.00,0.73]	[0.10, 0.85]	[0.54, 0.93]
Panel C: $\gamma_t^{expected}$									
γ_B/γ_G				0.47	0.47	0.47	0.24	0.24	0.24
Minimum ψ_G				0	0	0	0	0	0
90% CI, ψ_G				[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
Minimum ψ_B				0.53*	0.76*	0.88*	0.76*	0.88*	0.94*
90% CI, ψ_B				[0.44, 0.60]	[0.72, 0.80]	[0.86, 0.90]	[0.62, 0.87]	[0.81, 0.93]	[0.91, 0.97]
Panel D: $\gamma_t^{implied}$									
γ_B/γ_G							0.79	0.79	0.79
Minimum ψ_G							0	0	0
90% CI, ψ_G							[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
Minimum ψ_B							0.21*	0.60*	0.80*
90% CI, ψ_B							[0.08, 0.31]	[0.54, 0.65]	[0.77, 0.83]

Table 7
International evidence

For 20 stock market indexes, Panel A of the table reports the results of pooled panel regressions of the differences in the conditional monthly price of risk during good times (OECD non-recession months) to that of an investor who invests for twelve months during recessions, starting **after** the stock market reached its low during the recession. A positive value means that the price of risk is higher in recessions. We first compute conditional prices of risk within each month and thereafter investigate the average conditional prices of risk in normal times versus in recessions. We include stock index fixed effects and cluster standard errors by index and date. Using the point estimates of risk prices, we report in Panel B the empirical lower bound estimates of

$$\psi_B = 1 - \frac{\gamma_B}{\gamma_G}\frac{1-\psi_G}{k}. \tag{41}$$

where γ_B is the estimate of the price per unit of market variance risk in recessions (bad times) and γ_G is the non-recession estimate. The panels report the ratio of the risk prices, the minimum excess volatility ratio in good times, ψ_G, and the minimum excess volatility ratio in bad times – the fraction of the total market variance in good and bad times that is unrelated to fundamentals. Given the ratio $\frac{\gamma_B}{\gamma_G}$ and a value of $k = 2$ (price of spanned SDF risk doubles in recessions), we find the lowest possible $\psi_G \geq 0$ such that $\psi_B \geq 0$ in Equation (41). At this minimum value of ψ_G, we then compute the minimum ψ_B.* indicates statistical significance at the 10% level.

	Sharpe ratio	$\gamma_t^{\text{realized}}$	$\gamma_t^{\text{expected}}$
Panel A: Cyclicality in risk prices			
recession − non-recession	0.17*	0.99	−0.41
t-stat	2.48	0.65	−1.26
Panel B: Required excess volatility			
γ_B/γ_G		1.08	0.88
Minimum ψ_G		0	0
Minimum ψ_B		0.46	0.56

Table 8

Consumption growth, financial and macroeconomic conditions, and the price of risk

This table reports the results of regressions:

$$\text{Price of risk}_t = \alpha + \beta \times \text{Indicator}_t + \epsilon_t \tag{42}$$

t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors. Data on financial risk indicator is at the monthly horizon and obtained from the Chicago Fed database. NFCI is the National Financial Conditions Index. According to the Chicago Fed, "Risk" captures volatility and funding risk in the financial sector. Credit captures credit conditions and leverage consists of debt and equity measures. High values of the variables are historically associated with tighter-than-average conditions in financial markets, i.e., bad times. CFNAI is the Chicago Fed National Activity Index, built to capture movements in economic expansions and contractions as well as periods of increasing and decreasing inflationary pressure. A low value of this variable is typically associated with economic contractions. The "Rec. prob." variable is the recessions probability of Chauvet and Piger (2008). Consumption is the St. Louis Fed monthly growth in personal consumption expenditures of non-durable and service goods, deflated with the CPI. * indicates statistical significance at the 10% level.

	Sharpe ratio			Price per unit of variance risk		
Indicator	SR_t^{realized}	SR_t^{expected}	SR_t^{option}	$\gamma_t^{\text{realized}}$	$\gamma_t^{\text{expected}}$	γ_t^{option}
Panel A: National financial condition indicators						
NFCI	−0.16*	−0.02*	0.02	−5.47*	−0.65*	−0.88*
t-stat	−2.86	−2.32	1.56	−3.02	−4.03	−2.30
Risk	−0.16*	−0.02*	0.02	−5.43*	−0.60*	−0.96*
t-stat	−2.98	−1.82	1.44	−2.98	−3.42	−2.67
Credit	−0.01	−0.02	0.02	−1.49	−0.49*	−1.11*
t-stat	−0.16	−1.62	1.05	−0.62	−2.08	−2.28
Leverage	−0.17*	−0.01	0.01	−6.21*	−0.34	−0.17
t-stat	−3.05	−0.65	1.04	−3.36	−0.97	−0.35
Non-fin. leverage	−0.08*	0.01	−0.02*	−1.78	0.18	−0.67*
t-stat	−1.82	0.58	−2.29	−1.04	0.70	−4.03
Panel B: Other macroeconomic variables						
Recession prob.	−0.46*	−0.07*	0.03	−22.71*	−2.18*	−1.50*
t-stat	−1.98	−1.89	1.10	−3.49	−2.96	−3.65
CFNAI	0.04	0.01*	−0.00	3.24	0.28*	0.10*
t-stat	0.82	1.65	−0.91	1.52	1.73	1.72
Consumption	5.24*	0.26	−0.02	352.12	8.41	−7.23
t-stat	1.90	0.64	−0.11	1.56	0.97	−1.35

Table 9
Further international evidence for macro variables

This table reports the results of panel regressions on the form:

$$\text{Risk price}_t^i = \alpha + \beta \times \text{macro variable}_{t+1,t+s}^i + \epsilon_q \tag{43}$$

where i denotes indexes for the up to 20 international stock market indexes in our sample and macro variable$_{t+1,t+s}^i$ is either the: (i) eight quarters growth in consumption, (ii) the contemporaneous value of the market dividend-to-price ratio, or (iii) eight months growth in industrial production. t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors. 'Standardized' rows report results where we standardize both the risk price and the macro variable within each country before we pool the data for the regression. 'Raw' rows report results where we pool data without standardizing. * indicates statistical significance at the 10% level.

	Sharpe ratio		Price per unit of variance risk	
	SR_t^{realized}	SR_t^{expected}	$\gamma_t^{\text{realized}}$	$\gamma_t^{\text{expected}}$
Panel A: Consumption growth				
Raw	0.14*	−0.10	73.43*	3.35
t-stat	3.61	−0.23	3.29	0.48
Standardized	0.17*	0.01	0.14*	0.05
t-stat	4.66	0.14	3.61	0.47
Panel B: Dividend-price ratio				
Raw	−4.46*	−0.34	−99.20*	−11.28*
t-stat	−2.35	−0.55	−2.07	−1.72
Standardized	−0.09*	0.03	−0.08*	−0.04
t-stat	−3.84	0.58	−3.12	−0.81
Panel C: Industrial production				
Raw	3.47*	0.22	58.50*	5.10*
t-stat	4.67	1.33	3.62	1.83
Standardized	0.15*	0.07	0.10*	0.08*
t-stat	4.43	1.36	3.62	1.78

Table 10
Utility statistics

This table reports quarterly estimates of utility parameters. $\rho_{m,c,t}$ is the correlation between log consumption growth and log stock market returns, $\rho_{m,w,t}$ is the correlation between the returns on the log aggregate wealth portfolio and log stock market returns, and $\sigma_{c,t}, \sigma_{w,t}, \sigma_{m,t}$ are volatilities of log consumption growth, log returns on the wealth portfolio, and log returns on the stock market respectively. We report historical statistics for NBER recessions and normal times (non recession quarters). For consumption, we use data from St. Louis Fed on real personal consumption of non-durable and service goods. We use S&P 500 returns from CRSP as a proxy for stock market return. We create a proxy for the return on aggregate wealth from the asset wealth, consumption, and labor income in Lettau and Ludvigson (2001). The data is quarterly and the sample spans Q3-1947 to Q3-2019. There are 241 quarters of normal times and 48 quarters of recessions.

	$\rho_{m,c,t}$	$\rho_{m,w,t}$	$\sigma_{c,t}$	$\sigma_{w,t}$	$\sigma_{m,t}$
NBER recessions	0.1428	0.8139	0.0055	0.0209	0.1126
Normal times	0.1771	0.7609	0.0042	0.0167	0.0663
NBER recessions/Normal times	0.8068	1.0697	1.3026	1.2515	1.6984

Table 11

Equity term structure model parameters and simulation results

All parameters and simulation results are in annual terms. Results are based on 10,000 simulations of 200 years of the model described in Section 6. For each simulation period, expected excess returns and volatilities are calculated using the model solutions and the current x_t, and these are used to compute ex ante Sharpe ratios and γ_t. Results report averages across simulations for each statistic.

Panel A: Calibrated Parameter Values

Parameter	$\overline{x}$	ϕ_x	g	r^f	σ_d	σ_z	σ_x
Value	0.70	0.85	0.03	0.02	0.10	0.20	0.125

Panel B: Simulation Results

Unconditional averages

Maturity n (years)	Excess Return	$\sigma_{n,t}$	Sharpe	$\gamma_{n,t}$
1	0.07	0.17	0.39	2.26
2	0.07	0.18	0.39	2.24
10	0.07	0.19	0.37	1.96
30	0.07	0.19	0.35	1.81

Sharpe ratio: Cyclicality results

Maturity n (years)	Very bad times (x_t top decile)	Bad times ($x_t > \overline{x}$)	Good times ($x_t < \overline{x}$)
1	0.45	0.43	0.35
2	0.45	0.43	0.35
10	0.44	0.41	0.32
30	0.43	0.40	0.30

$\gamma_{n,t}$: Cyclicality results

Maturity n (years)	Very bad times (x_t top decile)	Bad times ($x_t > \overline{x}$)	Good times ($x_t < \overline{x}$)
1	1.90	2.14	2.37
2	1.89	2.14	2.36
10	1.77	1.94	1.98
30	1.69	1.83	1.79

Table 12
Excess volatility and investments

This table reports regressions on the form:

$$\Delta\text{Investment}_{q,q+s} = \alpha_s + \beta_s \times 1_{\gamma_q < 25\% \text{ quantile}} + \text{controls}_q + \epsilon_{q,q+s} \quad (44)$$

where $\Delta\text{Investment}_{q,q+s}$ is the relative change in investments from the end of quarter q to the end of quarter $q+s$. Data on investments is from the St. Louis FRED and we use the "Real Gross Private Domestic Investment" dataset. The dummy variable $1_{\gamma_q < 25\% \text{ quantile}}$ is 1 if, at the end of quarter q, the estimated (expected) price per unit of market variance is below its time-series 25% quantile — our proxy of high excess volatility periods. t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors. * indicates statistical significance at the 10% level.

Horizon quarters	$s=1$	$s=2$	$s=3$	$s=4$	$s=5$	$s=6$	$s=7$	$s=8$
Panel A: Univariate regression								
$1_{\gamma_t < 25\% \text{ quantile}}$	−0.01*	−0.02*	−0.03*	−0.03	−0.04	−0.03	−0.04	−0.04
t-stat	−2.04	−2.19	−1.72	−1.26	−1.05	−0.81	−0.80	−0.68
Panel B: Including controls								
$1_{\gamma_t < 25\% \text{ quantile}}$	−0.02*	−0.02*	−0.02	−0.03	−0.03	−0.03	−0.04	−0.03
t-stat	−2.14	−1.75	−1.06	−0.71	−0.76	−0.71	−0.97	−0.78
Controls:								
VIX	0.00*	0.00*	0.00*	0.00*	0.01*	0.01*	0.01*	0.01*
t-stat	2.23	3.10	3.05	3.20	3.36	3.95	3.79	2.87
Recession probabilities	0.04	0.03	0.02	0.03	0.06	0.06	0.07	0.05
t-stat	0.73	0.35	0.17	0.33	0.89	0.83	1.00	0.87
NFCI	−0.06*	−0.09*	−0.12*	−0.15*	−0.18*	−0.20*	−0.20*	−0.21*
t-stat	−2.48	−2.96	−3.07	−2.49	−2.26	−2.23	−3.07	−2.97
Dividend-price ratio	0.26	1.47	3.03	4.88	6.46	8.30*	9.06*	10.70*
t-stat	0.34	0.76	1.02	1.29	1.62	2.13	2.86	3.24
Credit spread	0.00	0.01	0.02	0.03	0.04	0.04	0.04*	0.04*
t-stat	0.41	0.67	0.88	1.13	1.41	1.58	2.09	2.16

Appendix

A. Details for Section 3

A.1. Details on option implied risk prices

As our third approach for computing the price per unit of variance risk, we look to option markets.[12] The premise for this approach is that the projection of the stochastic discount factor onto the market return, $M_{t+1}|R_{m,t+1} = \delta_t R_{m,t+1}^{\gamma_t}$, prices the market and derivatives written on the market. This premise is common in previous option literature, see e.g. Bliss and Panigirtzoglou (2004). Under this premise, we can relate the state price density ($\pi_{m,t+1}(x)$) of market returns to the physical probability density ($p_{m,t+1}(x)$) and a risk adjustment in the following way:

$$\pi_{m,t+1}(x) = p_{m,t+1}(x)\delta_t x^{-\gamma_t} \tag{45}$$

Using insights from Breeden and Litzenberger (1978), we can use option prices to back out risk-neutral densities, say $f_t^*(R_{m,t+1}) = \pi_{m,t+1}(x)R_t^f$. These densities reflect the time t real-time risk-adjusted probabilities over the potential future market outcomes.

Now, from Equation (45), we can write the stock market's physical probability distribution function, say $F_{m,t}(x)$, as

$$F_{m,t+1}(x) = \int_{-\infty}^{x} p_{m,t+1}(y)dy = \int_{-\infty}^{x} \frac{\pi_{m,t+1}(y)y^{-\gamma_t}}{\delta_t}dy \tag{46}$$

If we knew the true values of the parameters δ_t and γ_t then we could directly infer the stock market probability distribution from the observable state price density. However, the true values of the parameters are not directly observable and we therefore have to come up with a way to infer them. To achieve this task and infer the true values of δ_t and γ_t, we follow Bliss and Panigirtzoglou (2004) and use the so-called Berkowitz test, cf. Berkowitz (2001). The idea behind the Berkowitz test is that, for the true values of δ_t and γ_t, the distribution of $u_{t+1} = F_{m,t}(R_{m,t+1})$ is uniform and the distribution $y_{t+1} = \Phi^{-1}(u_{t+1})$ is standard normal. Therefore, to conduct the Berkowitz test, we estimate the coefficients in the regression model:

$$y_{t+1} = a + \beta y_t + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim N(0, \sigma^2) \tag{47}$$

and perform a likelihood ratio test of the joint hypothesis that $a = \beta = 0$ and $\sigma^2 = 1$.[13] It is worth noticing that, even though there might be momentum effects in returns, then we will still want $b = 0$ because the

[12] The empirical methodology in this section was first reported in Chapter 3 of Jensen (2018), which this paper now supersedes.

[13] We use non-overlapping monthly horizon distributions and returns. The hypothesis that $b = 0$ is therefore natural. For the case with overlapping returns see e.g. Bliss and Panigirtzoglou (2004) for a thorough discussion of the test.

true distribution should take these momentum effects into account. The Berkowitz likelihood ratio test for non-overlapping returns is:

$$LR = -2(LL(0,0,1) - LL(a,\beta,\sigma^2)) \sim \chi^2_3 \tag{48}$$

where $LL(a,\beta,\sigma^2))$ is the log likelihood of Equation (47). The likelihood ratio test statistic, LR, is chi-square distributed with three degrees of freedom.

To find the values of δ_t and γ_t, we minimize the Berkowitz test statistic in Equation (48) under the constraint that, for all dates t, the equation $\int_{-\infty}^{\infty} \frac{\pi_t(y)y^{-\gamma_{m,t}}}{\delta_{m,t}}dy = 1$ must hold. This constraint ensures that the resulting physical return distributions integrate to one at all points in time. Written in mathematical terms, the optimization problem is:

$$\min_{\delta_t} -2\left(LL(0,0,1) - LL(a,\beta,\sigma^2)\right) \tag{49}$$

$$s.t. \quad \gamma_t \text{ solves } \int_{-\infty}^{\infty} \frac{\pi_t(y)y^{-\gamma_t}}{\delta_t}dy = 1, \quad \text{for all } t \tag{50}$$

For a given value of δ_t, the constraints provide enough equations to solve for the time-varying γ_t. Specifically, for a given level of δ_t, at any point in time, we only have to solve for γ_t. If γ_t was linear in the constraint, then solving for the parameter would be straightforward. However, γ_t enters non-linearly in the constraint and we need to address this non-linearity. The generalized recovery methodology of Jensen, Lando, and Pedersen (2019) provides us with the argument we need. If we assume that there is a solution to the constraint for one given γ_t, then that solution is almost surely unique. Practically, this means that there will be at most one solution to the constraint equation.[14]

To optimize over the parameter, δ_t, we need to make an assumption on its functional form. We allow δ_t to be time varying through the time variation in the gross risk-free rate. Specifically, we assume that $\delta_t = \frac{1}{R_t^f} + c$ where R_t^f is the time t gross risk-free rate and c is a time-invariant parameter. This functional form of δ_t is conveniently simple while nesting the risk-neutral distribution as a solution if $c = 0$. To minimize Equation (49), we search over a grid of values for c and pick the c which provides the lowest Berkowitz test statistic. Importantly, for each value of c, the constraints ensure that we can infer a time varying level of γ_t. For different values of c, the γ_t time series will differ and consequently also the physical distributions, which gives us the variation in the Berkowitz test statistics that we need for our optimization.

We set γ_t^{option} to be the optimized γ_t from the optimization problem in 49. That is, γ_t^{option} takes the values that best reconciles the ex ante observable option prices with the ex post realized returns.

[14] Our methodology is closely related to the methodology used in Bliss and Panigirtzoglou (2004). In short, the difference in the two methodologies is that, they optimize over a constant level of stock market γ whereas we optimize over the stock market time preference parameter δ_t and allow γ_t to be time-varying through the constraints in Equation (49).

A.2. Inferring risk-neutral distributions from options

We compute risk-neutral distributions using a spline approach as suggested in Figlewski (2018). On each last trading day of the month, we compute the spline that best fits the observed volatility surface under the two conditions that: (i) the left part of the surface is monotonically decreasing (options with moneyness less than 1) and (ii) the resulting risk-neutral distribution is non-negative. Given a fitted spline, we compute the risk-neutral distribution as the second derivative of the resulting Black-Scholes prices. When we have option data with exactly one month maturity, we simply compute the distribution at this horizon. When we do not have such horizons, we compute distributions for the closest days (from above and below) for which there is data. We then linearly interpolate the distributions for these horizons to obtain a monthly horizon distribution.

There are of course many other ways in which we can compute risk-neutral distributions. For example, we can use a parametric approach like in Bates (2000), the ”fast-and-stable” method of Jackwerth (2004), or a number of alternative ways to essentially just smooth implied volatilities to obtain a smooth continuous price function that we can numerically differentiate. Since the Berkowitz (2001) test relies heavily on the first and second moments of the recovered risk-neutral distributions, which are almost identical using either of these methods, we are confident that alternative methods will yield similar cyclicalities in option implied prices of market variance risk when using the approach described above.

B. Details on option-implied risk prices by horizon

Next, we turn to option markets and show that the term structure of returns to particularly interesting option portfolios is directly related to the cyclicality in the price per unit of variance risk. We start by going through our theoretical setting and we present empirical evidence thereafter.

Theory

To simplify exposition, we consider an economy in which returns and the SDF are conditionally jointly log-normal, with

$$r_{m,t+1} \equiv \log R_{m,t+1} = \mu_{R,t} + \sigma_{\varepsilon,t}\varepsilon_{t+1} - \frac{1}{2}\sigma^2_{\varepsilon,t}, \tag{51}$$

$$m_{t+1} \equiv \log M_{t+1} = -r_{f,t} - \gamma_t r_{m,t+1} + \sigma_{\eta,t}\eta_{t+1} - \left[\frac{1}{2}(\sigma^2_{\eta,t} + \gamma_t(1+\gamma_t)\sigma^2_{\varepsilon,t}) - \gamma_t\mu_{R,t}\right], \tag{52}$$

where ε_{t+1} and η_{t+1} are standard normal and independent (both over time and with respect to each other). The last terms in both lines are Jensen's inequality corrections to ensure that $\log \mathbb{E}_t[R_{m,t+1}] = \mu_{R,t}$ and $\log \mathbb{E}_t[M_{t+1}] = -r_{f,t}$. Writing the unexpected part of m_{t+1} as a linear combination of a projection onto the market and an orthogonal term is without loss of generality in this setting. Define $\sigma_t^2 = \text{Var}_t(r_{m,t+1})$. Since $\mu_{R,t} - r_{f,t} = -\text{Cov}_t(m_{t+1}, r_{m,t+1})$, it follows that $(\mu_{R,t} - r_{f,t})/\sigma_t^2 = \gamma_t$. This motivates our use of γ_t to refer to the loading of the log SDF onto the market.

We now consider options. No arbitrage implies the existence of a risk-neutral density $f_t^*(R_{m,t+1})$ such that

$$f_t^*(R_{m,t+1}) = R_{f,t}\mathbb{E}_t[M_{t+1}|R_{m,t+1}]f_t(R_{m,t+1}), \tag{53}$$

where $f_t(R_{m,t+1})$ is the objective physical density. As observed by Schreindorfer and Sichert (2023), $f_t^*(R_{m,t+1})/R_{f,t}$ can be thought of as the price of the Arrow-Debreu security that pays 1 if the return $R_{m,t+1}$ is realized and 0 otherwise, while $f_t(R_{m,t+1})$ can be thought of as its expected payoff.[15] This implies that the log expected return of the Arrow-Debreu security is $-\log \mathbb{E}_t[M_{t+1}|R_{m,t+1}]$. Using the characterization in (51)–(52) and the assumption of log-normality, this implies that the log expected excess return is equal to

$$-\log \mathbb{E}_t[M_{t+1}|R_{m,t+1}] - r_{f,t} = \gamma_t r_{m,t+1} - \gamma_t \left(\mu_{R,t} - \frac{1}{2}\sigma_{\varepsilon,t}^2\right). \tag{54}$$

Now consider a strategy that goes short one unit of the AD security for return state $R_{m,t+1} = \omega_1$, and long one unit of the AD security for state $R_{m,t+1} = \omega_2$, where $\omega_2 > \omega_1$. This strategy can be thought of as a conditional binary bet: it involves a payoff of 1 if $R_{m,t+1} = \omega_2$, a payoff of -1 if $R_{m,t+1} = \omega_1$, and 0 otherwise. Denote the return on this strategy by $R_{\omega,t+1}$. Using (54), the log expected return on the strategy is

$$\begin{aligned} \log \mathbb{E}_t[R_{\omega,t+1}] &= \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_1] - \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_2] \\ &= \gamma_t(\omega_2 - \omega_1). \end{aligned} \tag{55}$$

Intuitively, γ_t is the market price per unit of variance risk. This strategy fixes the quantity of risk: it pays off either 1 or -1, and importantly, it holds fixed the return outcomes ω_1 and ω_2. For example, if $\omega_1 = 0\%$ and $\omega_2 = 2\%$, the bet is always over a 2-percentage-point range over the index value as of $t+1$ regardless of what the value of σ_t^2 is. The expected return therefore depends only on the price of risk γ_t.

Equation (55) characterizes the one-period (or short-horizon) log expected return on the above option strategy. We now consider the two-period (long-horizon) return on this strategy, in order to characterize the term structure of expected returns. To maintain notation, continue to set the option expiration date to $T = t+1$, but now step back to period $t-1$, two periods from expiration. The strategy's log expected return as of tomorrow (date t) will, as in (55), be higher given a positive shock to γ_t. If $\text{Cov}_{t-1}(\gamma_t, M_t) < 0$ — that is, if the price per unit of risk is higher in good times, as we found empirical evidence for in the previous

[15] This is loose only insofar as these are continuous densities. To formalize this fully, one can either consider a discretized version of the state space (as we will do in the empirical analysis) or define the Arrow-Debreu (AD) payoff to be a Dirac delta function.

section — then this implies that the return on the strategy from $t-1$ to t will be positive in bad times.[16] The strategy thus provides a hedge against shocks to M_t, implying that its two-period expected return should be lower than its one-period expected return.[17] This implies a downward-sloping term structure of risk prices for option portfolios that fix the quantity of risk in the sense described above. We now turn to index options data for suggestive empirical evidence along these lines.

Evidence[18]

We now seek to estimate returns on the fixed-quantity-of-risk strategy $\mathbb{E}_t[R_{\omega,T}]$, varying the horizon $T-t$. Our test of interest is whether this average return decreases in the horizon $T-t$, which (following the discussion above) would provide indirect ex ante evidence that γ_t is higher in good times.

As in the first line of (55), estimating $\mathbb{E}_t[R_{\omega,t+1}]$ is equivalent to estimating the ratio $\frac{\mathbb{E}_t[M_T|R_{m,T}=\omega_1]}{\mathbb{E}_t[M_T|R_{m,T}=\omega_2]}$. This ratio of SDF realizations across return states, holding fixed the difference in returns, can be estimated using risk-neutral probabilities $f_t^*(\omega) \equiv f_t^*(R_{m,0\to T}=\omega)$ obtained from index option prices. In particular, we use the same Breeden and Litzenberger (1978)–based approach as described in Section 3 to back out a discretized distribution $f_t^*(\omega)$ across possible returns ω realized over the life of the option. We then translate these into a set of conditional probabilities over binary outcomes — in particular, the probability that the index return from 0 to T will be ω_1 conditional on it being either ω_1 or ω_2. To estimate $\frac{\mathbb{E}_t[M_T|R_{m,T}=\omega_1]}{\mathbb{E}_t[M_T|R_{m,T}=\omega_2]}$, we then use the fact that

$$\frac{\pi_t^*}{1-\pi_t^*} = \phi_{t,T}\,\frac{\pi_t}{1-\pi_t}, \tag{56}$$

$$\text{where}\;\; \pi_t^* \equiv f_t^*(R_{m,0\to T}=\omega_1 \mid R_{m,0\to T}\in\{\omega_1,\omega_2\}),$$

$$\pi_t \equiv f_t(R_{m,0\to T}=\omega_1 \mid R_{m,0\to T}\in\{\omega_1,\omega_2\}),$$

$$\phi_{t,T} \equiv \frac{\mathbb{E}_t[M_T|R_{m,T}=\omega_1]}{\mathbb{E}_t[M_T|R_{m,T}=\omega_2]}.$$

While we can measure π_t^* from index options data directly, we must estimate π_t across horizons by using the

[16]Note that the unexpected return on the strategy from $t-1$ to t depends on not just the expected return from t to $t+1$, but also on $\log f_t(\omega_2) - \log f_t(\omega_1)$: the unexpected log return, from (53), depends on $\log f_t^*(\omega_2) - \log f_t^*(\omega_1) = \log \mathbb{E}_t[M_{t+1}|R_{m,t+1}=\omega_2] - \log \mathbb{E}_t[M_{t+1}|R_{m,t+1}=\omega_1] + \log f_t(\omega_2) - \log f_t(\omega_1)$. One concern might be that $\log f_t(\omega_2) - \log f_t(\omega_1)$ decreases in bad times enough to in fact make the strategy have a negative unexpected return in these times. But as shown in a later subsection (Section B.1), one can guarantee that $\log f_t(\omega_2) - \log f_t(\omega_1)$ increases in bad times (thereby guaranteeing that the unexpected return is higher in these times) by focusing on sufficiently high return states ω_1 and ω_2. (And more generally, the change in $\log f_t(\omega_2) - \log f_t(\omega_1)$ is likely to be quite small in practice.) So this is not, in our view, a first-order concern.

[17]A full formal analysis of the two-period expected return would require fully specifying the dynamics of all the state variables γ_t, σ_t^2, and so on.

[18]The empirical results described here were originally reported in Lazarus (2022), which this paper now supersedes.

fact that it must be an unbiased forecast of the terminal outcome $\mathbb{1}(R_{m,0\to T} = \omega_1 \mid R_{m,0\to T} \in \{\omega_1, \omega_2\})$ by definition. That is, we are effectively estimating the price of risk embedded in the π_t^* values across horizon such that the implied π_t values have zero average forecast error for the terminal index outcome. We provide formal details of this approach in the appendix (see Section B.2).

For implementation, we use S&P 500 index options data from the OptionMetrics database for the period 1996–2018. This yields data for 5,537 trading dates and 991 expiration dates. We drop any options with bid prices of zero (or less than zero), with Black-Scholes implied volatility of greater than 100 percent, or with greater than 12 weeks to maturity (given the relative lack of observations and statistical power beyond this maturity), and calculate each option's end-of-day price as the midpoint between its bid and ask prices.

For each observed expiration date T and associated initial option trading date 0, we define the relevant (sub)set of possible terminal index returns as

$$\Omega = R^f_{0,T} \exp\Big(\big\{[-0.10, -0.08), [-0.08, -0.06), \ldots, [0.06, 0.08), [0.08, 0.10)\big\}\Big).$$

In words, state ω_1 is said to be realized when the gross index-price appreciation, in excess of the risk-free rate $R^f_{0,T}$, is between $\exp(-0.1)$ and $\exp(-0.08)$, or equivalently when the log excess return is between -10% and -8%, and analogously for ω_2, and so on. We exclude all terminal states more than 10% out of the money (where moneyness is relative to a zero excess return) in either direction. Note that the states are equally spaced, and all binary bets (e.g., ω_2 vs. ω_3, or ω_5 vs. ω_6) have the same fixed 2-percentage point range of return outcomes within a given option contract, as required by construction. For a given option contract (i.e., a given set of option prices observed from 0 to T), we consider only (ω_i, ω_{i+1}) pairs for which the realized index return was either ω_i or ω_{i+1}. (Without this conditioning, the conditional physical probabilities would be undefined.) This leaves 549 observations (tuples (t, T, i)) at the one-day horizon, which declines monotonically to 222 observations at the 60-day horizon (equivalently, the 12-week horizon), which motivates our focus on 1- to 12-week horizons.

We present the option-implied prices of risk by horizon $\kappa = T - t$ in Figure 3. As the graph shows, the estimated price of risk is significantly downward-sloping as one increases the horizon κ. In other words, one needs a lower price of risk to rationalize the returns on fixed-quantity-of-risk bets at longer horizons. Equivalently, these bets have lower expected returns at longer horizons. As discussed at the end of the previous subsection, this implies that $\text{Cov}_{t-1}(\gamma_t, M_t) < 0$, so that the price per unit of risk is higher given good shocks. This is required in order for the strategy to provide a hedge against shocks to M_t and have lower expected returns at longer horizons.

These results are, in effect, an out-of-sample test in support of the procyclicality of γ_t established earlier in the paper. As discussed theoretically above, the option-implied term structure considered here is downward-sloping if and only if γ_t is procyclical ex ante. The fact that we indeed find a downward-sloping term structure therefore provides further support to the preceding results based on ex post returns on the market.

That said, due to the relatively short horizon necessitated by our data sample and cleaning, the evidence obtained from this estimation is at most suggestive: there is a downward-sloping term structure of implied expected returns on the above option strategy over a matter of weeks, implying that γ_t is procyclical at least at a weekly to monthly data frequency. Whether this speaks to slightly lower frequencies of data aggregation remains an interesting question for future work.

B.1. Further Details for Section B

This section continues the discussion in footnote 16 on sufficient conditions to guarantee that the unexpected return on the fixed-quantity-of-risk strategy is higher in bad times. As in that footnote, the unexpected log return on the strategy from $t-1$ to t depends on $\log f_t^*(\omega_2) - \log f_t^*(\omega_1) = \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_2] - \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_1] + \log f_t(\omega_2) - \log f_t(\omega_1)$ (using (53)). The log-normal density assumption gives that

$$\begin{aligned} &\log f_t(\omega_2) - \log f_t(\omega_1) \\ &\quad = -\log(\omega_2) + \log(\omega_1) - \frac{\left(\log(\omega_2) - \mu_{R,t} + \frac{1}{2}\sigma_t^2\right)^2 - \left(\log(\omega_1) - \mu_{R,t} + \frac{1}{2}\sigma_t^2\right)^2}{2\sigma_t^2}. \end{aligned}$$

This decreases in $\mu_{R,t}$ and it may either increase or decrease in σ_t^2, so one concern (as raised in the footnote) might be that $\log f_t(\omega_2) - \log f_t(\omega_1)$ decreases in bad times enough to in fact make the strategy have a negative unexpected return in these times. But since γ_t decreases in bad times, we must have that $d\gamma_t \propto d\mu_t - \gamma_t d(\sigma_t^2) < 0$ in these times, or $d\mu_t / d(\sigma_t^2) < \gamma_t$. So in order for $\log f_t(\omega_2) - \log f_t(\omega_1)$ to increase in bad times so that the unexpected return is guaranteed to be positive, one can see (after some algebra) that it is sufficient to have

$$\frac{\log(\omega_1) + \log(\omega_2)}{2} \frac{1}{\sigma_t^2} > \gamma_t - d\mu_t / d(\sigma_t^2) > 0.$$

One can always find large enough return states to guarantee that this is the case, meaning that the unexpected return will always be higher in bad times as long as we're focusing on sufficiently high ω_1 and ω_2.

B.2. Details on implementation

We begin with equation (56). We are interested in how the price of risk $\phi_{t,T}$ changes on average with the horizon $T-t$, for multiple possible return state pairs (ω_1, ω_2). We therefore assume that for arbitrary pairs of return states (ω_1, ω_2) and (ω_3, ω_4), if $\omega_2/\omega_1 = \omega_4/\omega_3$, then the associated ϕ values are equivalent (i.e., $\phi_{t,T,\omega_1,\omega_2} = \phi_{t,T,\omega_3,\omega_4}$). This is in effect an assumption of scale independence (as would hold under, e.g., CRRA preferences), since we will use a set of equally spaced return states for empirical implementation.

Second, we assume that $\phi_{t,T} = \phi_{T-t}$, so that the SDF ratio depends only on the horizon $T - t$. Both assumptions are in effect for the purposes of notational simplification so that we may pool estimates across return-state pairs and expiration dates below.

To derive moment conditions for estimation, we begin by rearranging (56) as

$$\pi_t = \frac{\pi_t^*}{\pi_t^* + \phi_{T-t}(1-\pi_t^*)}. \tag{57}$$

This equation says how the risk-neutral probability and ϕ_{T-t} together pin down the (unobserved) physical probability, which by definition must be unbiased: $\pi_t = \mathbb{E}_t[\mathbb{1}\{R_{m,0\to T} = \omega_1\} \mid R_{m,0\to T} \in \{\omega_1, \omega_2\}]$. Using this unbiasedness property,

$$\mathbb{E}_t\left[\mathbb{1}\{R_{m,0\to T} = \omega_1\} - \frac{\pi_t^*}{\pi_t^* + \phi_{T-t}(1-\pi_t^*)} \,\middle|\, R_{m,0\to T} \in \{\omega_1, \omega_2\}\right] = 0. \tag{58}$$

Note that the random variable $\mathbb{1}\{R_{m,0\to T} = \omega_1\}$ is observable as of date T, as it simply indexes whether the terminal index return is equal to ω_1. Thus every value in (58) is in principle observable aside from π_{T-t}, so applying the law of iterated expectations to this equation yields a nonlinear moment condition for ϕ_{T-t} that can be estimated using the generalized method of moments (GMM).

One possible concern with such estimation is the likelihood of price measurement error affecting the measured risk-neutral probabilities in (58) given, for example, market microstructure noise. The GMM framework here, however, allows us to account for this noise without needing to estimate its magnitude separately. To discuss this estimation, we first generalize the notation slightly, and allow for arbitrary return states indexed by j, (ω_j, ω_{j+1}). We then assume that the observed conditional risk-neutral belief $\widehat{\pi}_{t,j}^*$ is measured with additive error with respect to the true value $\pi_{t,j}^*$ used in (58):

$$\widehat{\pi}_{t,j}^* = \pi_{t,j}^* + \epsilon_{t,j}, \tag{59}$$

where $\mathbb{E}[\epsilon_{t+k,j}\, \pi_{t+k',j}^* \mid R_{m,0\to T} \in \{\omega_j, \omega_{j+1}\}] = 0$ for all k, k', and $\epsilon_{t,j}$ follows an MA(q) for some value q. Using this and then Taylor expanding the observed analogue for the second term in (58), we obtain

$$\frac{\widehat{\pi}_{t,j}^*}{\widehat{\pi}_{t,j}^* + \phi_{T-t}(1-\widehat{\pi}_{t,j}^*)} = \frac{\pi_{t,j}^*}{\pi_{t,j}^* + \phi_{T-t}(1-\pi_{t,j}^*)} + \epsilon_{t,j} + \underbrace{\mathcal{O}\Big(\big(\epsilon_{t,j} + (\phi_{T-t} - 1)\big)^2\Big)}_{\text{higher-order terms}} \tag{60}$$

as $\epsilon_{t,j} \to 0$ and $\phi_{T-t} \to 1$,[19] where the latter limit $\phi_{T-t} = 1$ corresponds to the case of risk-neutrality as seen in (56).

[19] More formally, one may write the remainder term as $\mathcal{O}((\|\epsilon_{t,j}\| + (\phi_{T-t} - 1))^2)$ as $\|\epsilon_{t,j}\| \to 0$ and $\phi_{T-t} \to 1$, where $\|\epsilon_{t,j}\|$ indexes the bounds on $\epsilon_{t,j}$.

Thus equation (58) can be re-expressed up to higher-order terms as

$$\mathbb{E}_t\left[\mathbb{1}\{R_{m,0\to T} = \omega_1\} - \frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1-\widehat{\pi}^*_{t,j})} \;\middle|\; R_{m,0\to T} \in \{\omega_j, \omega_{j+1}\} \right] = -\epsilon_{t,j}. \tag{61}$$

The risk-neutral probabilities used on the left side of this equation are now the observable values (inclusive of noise, unlike the ideal values used in (56)). Since $\epsilon_{t,j}$ is assumed to follow an MA(q) as discussed above, we can then form a set of unconditional moments by instrumenting using lagged values of $\widehat{\pi}^*_{t,j}$, for any lags greater than q.

That is, defining the N-dimensional instrument vector

$$Z_{t,j} \equiv \begin{pmatrix} \widehat{\pi}^*_{t-q-1,j} \\ \vdots \\ \widehat{\pi}^*_{t-\overline{q},j} \end{pmatrix} \tag{62}$$

for some $\overline{q} > q$, we can then obtain the time-unconditional orthogonality condition

$$\mathbb{E}\left[\left(\mathbb{1}\{R_{m,0\to T} = \omega_j\} - \frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1-\widehat{\pi}^*_{t,j})} \mathbb{1}\Big\{ R_{m,0\to T} \in \{\omega_j, \omega_{j+1}\} \Big\} \right) Z_{t,j} \right] = 0. \tag{63}$$

This unconditional moment restriction is now amenable to empirical estimation over many expiration dates T, horizons $T-t$, and state pairs j. One can set the sample version of (63) to zero over all pairs $t = \tau_1, T = \tau_2$ such that $\tau_2 - \tau_1 = \kappa$, in order to identify ϕ_κ. One can then stack the moment condition for values of $\kappa = 1, 2, \ldots$, to obtain horizon-dependent risk-price estimates.

For empirical estimation, we define the set of return states (and, by implication, return-state pairs) Ω as discussed in the main text. (A given return state is realized if the excess log return is in a given 2-ppt range.) We use the S&P index options data set (and associated data filters) described in the main text, and we estimate risk-neutral densities as described in the appendix of Lazarus (2022) (see Section 3 for an intuitive discussion of such risk-neutral estimation). We restrict ϕ_{T-t} to be fixed by weeks to expiration, so ϕ_1 is, e.g., the one-week-horizon estimated risk price. Finally, we use the five-day-lagged observed risk-neutral probability $\widehat{\pi}^*_{t-5,j}$ as an instrument in the moment equation for $\widehat{\pi}^*_{t,j}$; following the discussion above, this is equivalent to assuming an MA(4) measurement-noise process and setting $\overline{q} = q + 1 = 5$.

Estimation for the ϕ_κ values shown in Figure 3 is then conducted with two-step GMM. The first-stage weight matrix is $Z'Z/\mathcal{T}$, where Z is the data matrix for the instruments and $\mathcal{T}$ is the number of observations. The second-stage weight matrix is then clustered by blocks of 8 time-adjacent observations. The figure presents the resulting estimates, which are downward-sloping by horizon; see the main text for additional discussion.